

The Optimality of the Interleaving Distance on Multidimensional Persistence Modules

Michael Lesnick*

Stanford University

July 15, 2011

Abstract

Building on an idea of Chazal et al. [11], we introduce and study the *interleaving distance*, a pseudometric on isomorphism classes of multidimensional persistence modules.

We present five main results about the interleaving distance. First, we show that in the case of ordinary persistence, the interleaving distance is equal to the bottleneck distance on tame persistence modules. Second, we prove a theorem which implies that the restriction of the interleaving distance to finitely presented multidimensional persistence modules is a metric. The same theorem, together with our first result, also yields a converse to the algebraic stability theorem of [11]; this answers a question posed in that paper. Third, we observe that the interleaving distance is stable in three senses analogous to those in which the bottleneck distance is known to be stable. Fourth, we introduce several notions of optimality of metrics on persistence modules and show that when the underlying field is \mathbb{Q} or a field of prime order, the interleaving distance is optimal with respect to one of these notions. This optimality result, which is new even for ordinary persistence, is the central result of the paper. We also prove that a version of this result holds for ordinary persistence modules over any field, provided we restrict attention to a class of well behaved ordinary persistence modules containing the finitely presented ones. Fifth, we show that the computation of the interleaving distance between two finitely presented multidimensional persistence modules M and N reduces to deciding the solvability of $O(\log m)$ systems of multivariate quadratic equations, each with $O(m^2)$ variables and $O(m^2)$ equations, where m is the total number of generators and relations in a minimal presentation for M and a minimal presentation for N .

*mlesnick@stanford.edu

1 Introduction

1.1 Background

Persistent homology [24, 33, 23], a multiscale extension of the classical homology functor of algebraic topology, has proven to be an extremely useful tool in applied and computational topology. In the last decade, it has been applied effectively to a wide variety of problems of practical interest [20, 6, 13, 30, 24, 3, 16] and has established itself as a technical pillar of the emerging field of topological data analysis [4]. At the same time, persistent homology has been the focus of a growing body of theoretical work. This work has offered insight into algebraic and algorithmic questions surrounding persistent homology [11, 5, 29], and has begun to put on firm mathematical footing the use of persistent homology in an inferential setting [17, 11, 12, 14, 13, 15].

Persistent homology enjoys a number of nice properties that make it a very attractive tool with which to work. First, the isomorphism classes of persistence modules, the algebraic targets of the persistent homology functors, are completely described by invariants called persistent diagrams which are readily visualized and which transparently reflect geometric properties of the source objects [33]. Persistence diagrams can be efficiently computed from geometric input [33].

Another favorable property of persistent homology, and one that will be of particular interest to us here, is that there are well-behaved, easily understood, and readily computable metrics on persistence modules. The most popular of these is known as the bottleneck distance. The bottleneck distance and its variants [18] play a central role in both the theory and applications of persistent homology: The stability theorems for persistence [17, 11, 12] and theorems about inferring persistent homology from point cloud data [14, 13] are typically formulated using bottleneck distance, and many applications of persistent homology to shape comparison and related tasks [31, 9, 3, 1, 18] rely in an essential way on computations of the bottleneck distance and its variants.

In 2006 the authors of [8] introduced multidimensional persistent homology, a generalization of persistent homology. Whereas ordinary persistent homology produces algebraic invariants of topological spaces filtered by a single real parameter, multidimensional persistence produces algebraic invariants of topological spaces filtered by several real parameters.

Such “multifiltered” geometric objects arise naturally in a number of settings of interest in applications; for example, as we discuss in Section 7.4, there are natural ways of defining functors which associate topological spaces filtered by n parameters to \mathbb{R}^n -valued functions on topological spaces, or topological spaces filtered by $n + 1$ parameters to \mathbb{R}^n -valued functions on metric spaces. By way of these functors, multidimensional persistence provides invariants of such functions that are capable of encoding far more geometric information than their 1-D persistence analogues.

In fact, this approach also allows us to construct rich families of invariants of metric spaces lacking the additional data of a function to begin with; to do this, we associate to each metric space X a function $f_X : X \rightarrow \mathbb{R}^n$ and then apply the multidimensional persistence invariants

available for functions on metric spaces. For example, as suggested in [8], f_X can be chosen to be a density estimator or an eccentricity function.

Despite the fact that in many settings multidimensional persistence yields a far richer set of invariants than ordinary persistence does, the computational topology community has thus far been slow to make use of the machinery of multidimensional persistence in applications. This is in sharp contrast to the case of ordinary persistence.

That the community has not been quicker to put multidimensional persistence to use is not entirely surprising. After all, with the added power and generality of multidimensional persistence comes a significant degree of additional mathematical complexity. In turn, this added complexity presents obstacles to extending in naive ways the usual methodologies for applying persistent homology.

For example, whereas the structure theorem for finitely generated ordinary persistence modules [33] gives that ordinary persistence modules decompose uniquely into cyclic summands, multidimensional persistence modules admit no such decomposition in general. As a result, the characterization of the isomorphism class of ordinary persistence modules in terms of persistence diagrams does not extend to the multidimensional case.

In the absence of an analogue of the persistence diagram for multidimensional persistence modules, the bottleneck distance between persistence modules does not admit a naive generalization to a metric in the multidimensional setting.

This being the case, the question of how to best generalize the bottleneck distance to the setting of multidimensional persistence has remained open.¹

In this paper, we address this question, motivated by the view that a sound theoretical understanding of how to best choose metrics for multidimensional persistence modules promises to facilitate the adaptation to the multidimensional setting of theoretical results and applications of ordinary persistence which require having a metric on persistence diagrams.

1.2 Overview

We introduce and study here the *interleaving distance*, a pseudometric on isomorphism classes of multidimensional persistence modules which restricts to a metric on finitely presented persistence modules. We define the interleaving distance in terms of ϵ -*interleaving homomorphisms*; these are generalizations to the setting of multidimensional persistence of objects introduced in the context of 1-D persistence in [11].² While the interleaving distance for ordinary persistence modules is not explicitly defined in [11], the definition is considered implicitly in the conclusion of that paper.

We present five main results about the interleaving distance. The first result, Theorem 5.2, shows that in the case of ordinary persistence, the interleaving distance is in fact equal to the bottleneck distance on tame persistence modules. Our proof relies on a generalization of the

¹The matter of defining a pseudometric on multidimensional persistence modules has, however, previously been considered in [10].

²What we call an ϵ -interleaving homomorphism here is called a **strong** ϵ -interleaving homomorphism in [11]; since we have no occasion to consider weak ϵ -interleaving homomorphisms here, we drop the descriptor strong.

structure theorem [33] for finitely generated ordinary persistence modules to (discrete) tame persistence modules. This generalization is proven e.g. in [32].

Our second main result is Theorem 6.1. As an immediate consequence of this theorem, we have Corollary 6.2, which says that the interleaving distance restricts to a metric on finitely presented persistence modules. Theorems 5.2 and 6.1 together also yield Corollary 6.3, a converse to the algebraic stability theorem of [11]. This result answers a question posed in [11].

Our third result is the observation that the interleaving distance is stable in three senses analogous to those in which the bottleneck distance is known to be stable. These stability results, while notable, require very little mathematical work; two of the stability results turn out to be trivial, and the third follows from a minor modification of an argument given in [12].

Our fourth main result, Corollary 10.2, is an optimality result for the interleaving distance. It tells us that when the underlying field is \mathbb{Q} or a field of prime order, the interleaving distance is stable in a sense analogous to that which the bottleneck distance is shown to be stable in [17, 11], and further, that the interleaving distance is, in a uniform sense, the most sensitive of all stable pseudometrics. This “maximum sensitivity” property of the interleaving distance is equivalent to the property that, with respect to the interleaving distance, multidimensional persistent homology preserves the metric on source objects as faithfully as is possible for any choice of stable pseudometric on multidimensional persistence modules; see Remark 9.1 for a precise statement. Our optimality result is new even for 1-D persistence. In that case, it offers some mathematical justification, complementary to that of [17, 11], for the use of the bottleneck distance.

In fact, provided we restrict attention to a class of well behaved ordinary persistence modules containing the finitely presented ones, the assumption that the underlying field is \mathbb{Q} or a field of prime order is unnecessary; our Theorem 10.6 gives an analogue of Corollary 10.2 for this class of modules, over arbitrary fields.

The main step in the proof of Corollary 10.2 is the proof of Theorem 10.7, which gives a condition equivalent to the existence of ϵ -interleaving homomorphisms between two persistence modules. Theorem 10.7 expresses transparently the sense in which ϵ -interleaved persistence modules are algebraically similar; we believe that this result is of independent interest.

Given our first four main results, it is natural to ask if and how the interleaving distance can be computed. Our fifth main result speaks to this question. The result, which follows from Theorem 11.4 and Proposition 11.7, is that the computation of the interleaving distance between two finitely presented multidimensional persistence modules M and N reduces to deciding the solvability of $O(\log m)$ systems of multivariate quadratic equations, each with $O(m^2)$ variables and $O(m^2)$ equations, where m is the total number of generators and relations in a minimal presentation for M and a minimal presentation for N . This result is just a first step towards understanding the problem of computing the interleaving distance; we plan to address the problem more fully in a subsequent paper.

1.2.1 Related Work

After making the first preprint of this paper publicly available, it was brought to our attention that the paper [19] contains an optimality result for the bottleneck distance similar to the optimality results given here, for the special case of 0-dimensional ordinary persistent homology. Our Theorem 10.6 generalizes a slight weakening of [19, Theorem 32]; see Remark 10.1.

1.3 Outline

We conclude this introduction with an outline of the rest of the paper. Sections 2-5 of the paper are primarily algebraic. Section 2 covers a variety of algebraic preliminaries that will be needed for the rest of the paper. In particular, we define multidimensional persistence modules as n -graded modules over a monoid ring $k[\mathbb{R}_{\geq 0}^n]$, define the interleaving distance, and discuss minimal presentations of persistence modules. Section 3 is devoted to the review of results about 1-D persistence which we use in Sections 4-6. The focus of Section 4 is on the adaptation of the structure theorem of [7] for discrete tame persistence modules to a class of well behaved persistence modules. Using this result, in Section 5 we prove Theorem 5.2, our first main result, which tells us that the interleaving distance is equal to the bottleneck distance for tame 1-D persistence modules. In Section 6, we prove our second main result, Theorem 6.1.

In Section 7, we review preliminaries of a geometric nature which we need in Sections 8-10. We define here three variants of the multidimensional persistent homology functor which are of interest from the standpoint of applications, and review CW homology and the stability of ordinary persistence. In Section 8 we present our third main result, that multidimensional persistent homology is stable with respect to the interleaving distance in three senses. In Section 9, we introduce a general framework for defining the optimality of pseudometrics on multidimensional persistence modules. We specialize this framework to arrive at several notions of optimality of such pseudometrics.

In Section 10, we prove Theorem 10.7, our characterization of the existence of ϵ -interleaving homomorphisms between two modules. Using this, we prove the optimality result Corollary 10.2 for the interleaving distance, our fourth main result. In addition, we present Theorem 10.6, our analogue of Corollary 10.2 for arbitrary fields in the case where we restrict attention to a class of well behaved 1-D persistence modules.

In Section 11, we present our fifth main result, the reduction of the computation of the interleaving distance between finitely presented persistence modules to the problem of deciding the solvability of systems of quadratics.

Section 12 concludes the paper with a discussion of open problems and future directions for research related to our investigations here.

2 Algebraic Preliminaries

In this section we define persistence modules and review some (primarily) algebraic facts and definitions which we will need throughout the paper.

2.1 First Definitions and Notation

2.1.1 Basic Notation

Let k be a field. Let \mathbb{N} denote the natural numbers, $\mathbb{Z}_{\geq 0}$ denote the non-negative integers, $\mathbb{R}_{\geq 0}$ denote the non-negative reals, and $\mathbb{R}_{> 0}$ denote the positive reals. We view \mathbb{R}^n as a partially ordered set, with $(a_1, \dots, a_n) \leq (b_1, \dots, b_n)$ iff $a_i \leq b_i$ for all i . Let \mathbf{e}_i denote the i^{th} standard basis vector in \mathbb{R}^n .

For $A \subset \mathbb{R}$ any subset, let \bar{A} denote $A \cup \{-\infty, \infty\}$.

2.1.2 Notation Related to Categories

For a category C , let $\text{obj}(C)$ denote the objects of C and let $\text{obj}^*(C)$ denote the set of isomorphism classes of objects of C . For $X, Y \in \text{obj}(C)$ let $\text{hom}(X, Y)$ denote the set of morphisms from X to Y .

2.1.3 Metrics, Pseudometrics, and Semi-pseudometrics

Recall that a *pseudometric* on a set X is a function $d : X \times X \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ with the following three properties:

1. $d(x, x) = 0$ for all $x \in X$.
2. $d(x, y) = d(y, x)$ for all $x, y \in X$.
3. $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$.

We'll often use the term *distance* in this paper as a synonym for pseudometric.

A *metric* is a pseudometric d with the additional property that $d(x, y) \neq 0$ whenever $x \neq y$. We define a *semi-pseudometric* to be a function $d : X \times X \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ satisfying properties 1 and 2 above.

2.1.4 Metrics on Categories

In this paper we'll often have the occasion to define a pseudometric on $\text{obj}^*(C)$, for C some category. For d such a pseudometric, $M, N \in \text{obj}(C)$, and $[M], [N]$ the isomorphism classes of M and N , we'll always write $d(M, N)$ as shorthand for $d([M], [N])$.

2.2 Commutative Monoids and Commutative Monoid Rings

Monoid rings are generalizations of polynomial rings.

A *commutative monoid* is a pair $(G, +_G)$, where G is a set and $+_G$ is an associative, commutative binary operation on G with an identity element. Abelian groups are by definition commutative monoids with the additional property that each element has an inverse. We'll often denote the monoid $(G, +_G)$ simply as G . A submonoid of a monoid is defined in the obvious way, as is an isomorphism between two monoids.

Given a set S , let $k[S]$ denote the vector space of formal linear combinations of elements of S . If $\bar{G} = (G, +_G)$ is a monoid, then the operation $+_G$ induces a ring structure on $k[G]$, where

multiplication is characterized by the property $(k_1 g_1)(k_2 g_2) = k_1 k_2 (g_1 +_G g_2)$ for $k_1, k_2 \in k$, $g_1, g_2 \in G$. We call the resulting ring the *monoid ring* generated by \bar{G} , and we denote it $k[\bar{G}]$.

Let A_n denote $k[x_1, \dots, x_n]$, the polynomial ring in n variables with coefficients in k . For $n > 0$, $\mathbb{Z}_{\geq 0}^n$ is a monoid under the usual addition of vectors. It's easy to see that $k[\mathbb{Z}_{\geq 0}^n]$ is isomorphic to A_n .

Similarly, $\mathbb{R}_{\geq 0}^n$ is a monoid under the usual addition of vectors. Let B_n denote the monoid ring $k[\mathbb{R}_{\geq 0}^n]$. We may think of B_n as an analogue of the usual polynomial ring in n -variables where exponents of the indeterminates are allowed to take on arbitrary non-negative real values rather than only non-negative integer values. With this interpretation in mind, we'll often write (r_1, \dots, r_n) as $x_1^{r_1} x_2^{r_2} \cdots x_n^{r_n}$, for $(r_1, \dots, r_n) \in \mathbb{R}_{\geq 0}^n$.

2.3 Multidimensional Persistence Modules

We first review the definition of a multidimensional persistence module given in [8]. We then define analogues of these over the ring B_n .

In what follows, we'll often refer to multidimensional persistence modules simply as “persistence modules.”

2.3.1 \mathbf{A}_n -Persistence Modules

Fix $n \in \mathbb{N}$; let \mathbf{A}_n denote the ring $k[x_1, \dots, x_n]$. Let \mathbf{e}_i denote the i^{th} standard basis vector in \mathbb{Z}^n . An **\mathbf{A}_n -persistence module** is an \mathbf{A}_n -module \mathbf{M} with a direct sum decomposition as a k -vector space $\mathbf{M} \cong \bigoplus_{u \in \mathbb{Z}^n} \mathbf{M}_u$ such that the action of \mathbf{A}_n on \mathbf{M} satisfies $x_i(\mathbf{M}_u) \subset \mathbf{M}_{u+\mathbf{e}_i}$ for all $u \in \mathbb{Z}^n$. In other words, a \mathbf{A}_n -persistence module is simply an \mathbf{A}_n -module endowed with an n -graded structure.

For \mathbf{M} and \mathbf{N} \mathbf{A}_n -persistence modules, we define $\text{hom}(\mathbf{M}, \mathbf{N})$ to consist of module homomorphisms $f : \mathbf{M} \rightarrow \mathbf{N}$ such that $f(\mathbf{M}_u) \subset \mathbf{N}_u$ for all $u \in \mathbb{Z}^n$. This defines a category whose objects are the \mathbf{A}_n -persistence modules. Let $\mathbf{A}_n\text{-mod}$ denote this category.

2.3.2 B_n -persistence modules

In close analogy with the definition of an n -graded \mathbf{A}_n -module, we define a **B_n -persistence module** to be a B_n -module M with a direct sum decomposition as a k -vector space $M \cong \bigoplus_{u \in \mathbb{R}^n} M_u$ such that the action of B_n on M satisfies $x_i^\alpha(M_u) \subset M_{u+\alpha\mathbf{e}_i}$ for all $u \in \mathbb{R}^n$, $\alpha \geq 0$.

For M and N B_n -persistence modules, we define $\text{hom}(M, N)$ to consist of module homomorphisms $f : M \rightarrow N$ such that $f(M_u) \subset N_u$ for all $u \in \mathbb{R}^n$. This defines a category whose objects are the B_n -persistence modules. Let $B_n\text{-mod}$ denote this category.

Our notational convention will be to use boldface to denote \mathbf{A}_n -persistence modules and italics to denote B_n -persistence modules. We'll often refer to \mathbf{A}_1 -persistence modules and B_1 -persistence modules as *ordinary* persistence modules.

2.3.3 On the Relationship Between \mathbf{A}_n -persistence Modules and B_n -persistence Modules

Since \mathbf{A}_n is a subring of B_n , we can view B_n as an \mathbf{A}_n -module. If \mathbf{M} is an \mathbf{A}_n -persistence module then $\mathbf{M} \otimes_{\mathbf{A}_n} B_n$ is a B_n -module. Further, $\mathbf{M} \otimes_{\mathbf{A}_n} B_n$ inherits an n -grading from those on B_n and \mathbf{M} which gives $\mathbf{M} \otimes_{\mathbf{A}_n} B_n$ the structure of a B_n -persistence module.

In fact, $(\cdot) \otimes_{\mathbf{A}_n} B_n$ defines a functor from $\mathbf{A}_n\text{-mod}$ to $B_n\text{-mod}$. It can be checked that this functor is fully faithful and descends to an injection on isomorphism classes of objects. Thus the functor induces an identification of $\mathbf{A}_n\text{-mod}$ with a subcategory of $B_n\text{-mod}$.

In light of this, we can think of B_n -persistence modules as generalizations of \mathbf{A}_n -persistence modules. Finitely presented B_n -persistence modules arise naturally in applications, as discussed in Section 7.4. In a sense that can be made precise using machinery mentioned in Remark 4.1, it is possible to view them as \mathbf{A}_n -persistence modules endowed with some additional data. However, this is awkward from the standpoint of constructing pseudometrics between B_n -persistence modules. We thus regard B_n -persistence modules as the fundamental objects of interest here, and use \mathbf{A}_n -persistence modules in this paper only in the case $n = 1$ to translate results about \mathbf{A}_1 -persistence modules into analogous results about B_1 -persistence modules.

In the remainder of Section 2.3, we present some basic definitions related to B_n -persistence modules. All of these definitions have obvious analogues for \mathbf{A}_n -persistence modules; we'll use these analogues where needed without further comment.

2.3.4 Homogeneity

Let M be a B_n -persistence module. For $u \in \mathbb{R}^n$, we say that M_u is a *homogeneous summand* of M . We refer to an element $v \in M_u$ as a *homogeneous element* of grade u , and write $gr(v) = u$. A *homogeneous submodule* of a B_n -persistence module is a submodule generated by a set of homogeneous elements. The quotient of a B_n -persistence module M by a homogeneous submodule of M is itself a B_n -persistence module; the n -graded structure on the quotient is induced by that of M .

2.3.5 Transition Maps

For M a B_n -persistence module, and any $u \leq v \in \mathbb{R}^n$, the restriction to M_u of the action on M of the monomial $x_1^{v_1-u_1} x_2^{v_2-u_2} \dots x_n^{v_n-u_n}$ defines a linear map with codomain M_v . Denote this map by $\varphi_M(u, v)$.

2.3.6 Shifts of B_n -Persistence Modules

Let M be a B_n -persistence module. For $u \in \mathbb{R}^n$, define $M(u)$ by taking, for all $v \in \mathbb{R}^n$, $M(u)_v = M_{u+v}$. We take the transition maps for $M(u)$ to be induced by those of M in the obvious way. Let $\vec{1} \in \mathbb{R}^n$ denote the vector whose components are each 1. As a matter of notational convenience, for $u \in \mathbb{R}^n$ and $\epsilon \in \mathbb{R}$, let $u + \epsilon$ denote $u + \epsilon \vec{1}$. For $\epsilon \in \mathbb{R}$, define $M(\epsilon)$ to be $M(\epsilon \vec{1})$. More generally, for any subset $Q \subset M$, let $Q(\epsilon) \subset M(\epsilon)$ denote the image of

Q under the bijection between M and $M(\epsilon)$ induced by the identification of each summand $M(\epsilon)_u$ with $M_{u+\epsilon}$.

Note that for any two B_n -persistence modules M and N , and $\epsilon \in \mathbb{R}_{\geq 0}$, a homomorphism $f : M \rightarrow N$ induces in an obvious way a homomorphism with domain $M(\epsilon)$ and codomain $N(\epsilon)$. By slight abuse of notation, we'll also refer to this induced map as f .

For a B_n -persistence module M and $\epsilon \in \mathbb{R}_{\geq 0}$ let $S(M, \epsilon) : M \rightarrow M(\epsilon)$, the **(diagonal) ϵ -shift homomorphism**, be the homomorphism whose restriction to M_u is the linear map $\varphi_M(u, u + \epsilon)$ for all $u \in \mathbb{R}^n$.

2.3.7 Tameness

Following [11] we'll call a B_n -persistence module *tame* if each homogeneous summand of the module is finite dimensional. Note that this is a more general notion of tameness than that which appears in the original paper on the stability of persistence [17].

2.4 ϵ -interleavings and the Interleaving Distance

We now define the interleaving distance on B_n -persistence modules.

For $\epsilon \geq 0$, we say that two B_n -persistence modules M and N are **ϵ -interleaved** if there exist homomorphisms $f : M \rightarrow N(\epsilon)$ and $g : N \rightarrow M(\epsilon)$ such that $g \circ f = S(M, 2\epsilon)$ and $f \circ g = S(N, 2\epsilon)$; we refer to such f and g as ϵ -interleaving homomorphisms.

The definition of ϵ -interleaving homomorphisms was introduced for B_1 -persistence modules in [11].

Remark 2.1. It's easy to see that if $0 \leq \epsilon_1 \leq \epsilon_2$ and M and N are ϵ_1 -interleaved, then M and N are ϵ_2 -interleaved.

We define $d_I : \text{obj}^*(B_n\text{-mod}) \times \text{obj}^*(B_n\text{-mod}) \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$, the **interleaving distance**, by taking $d_I(M, N) = \inf\{\epsilon \in \mathbb{R}_{\geq 0} \mid M \text{ and } N \text{ are } \epsilon\text{-interleaved}\}$.

Note that d_I is pseudometric. However, the following example shows that d_I is not a metric.

Example 2.1. Let M be the B_1 -persistence module with $M_0 = k$ and $M_a = 0$ if $a \neq 0$. Let N be the trivial B_1 -persistence module. Then M and N are not isomorphic, and so are not 0-interleaved, but it is easy to check that M and N are ϵ -interleaved for any $\epsilon > 0$. Thus $d_I(M, N) = 0$.

2.5 Free B_n -persistence Modules and Related Algebraic Basics

2.5.1 n -graded Sets

We begin our discussion of free B_n -persistence modules with some foundational definitions.

Define an n -graded set to be a pair $G = (\bar{G}, \iota_G)$ where \bar{G} is a set and $\iota_G : G \rightarrow \mathbb{R}^n$ is any function. When ι_G is clear from context, as it will usually be, we'll write $\iota_G(y)$ as $gr(y)$ for $y \in \bar{G}$. We'll sometimes abuse notation and write G to mean the the set \bar{G} when no confusion is likely. The union of disjoint graded sets is defined in the obvious way. For $\epsilon \geq 0$ and $G = (\bar{G}, \iota_G)$ an n -graded set, let $G(\epsilon)$ be the n -graded set (\bar{G}, ι'_G) , where $\iota'_G(y) = \iota(y) - \epsilon$.

For G an n -graded set, define $gr(G) : \mathbb{R}^n \rightarrow \mathbb{Z}_{\geq 0} \cup \{\infty\}$ by taking $gr(G)(u)$ to be the number of elements $y \in G$ such that $gr(y) = u$. Note that for any B_n -persistence module M , a set Y of homogeneous elements of M inherits the structure of an n -graded set from the graded structure on M , so that $gr(Y)$ is well defined.

2.5.2 Free B_n -persistence modules

The usual notion of a free module extends to the setting of B_n -persistence modules as follows: For G an n -graded set, let $\langle G \rangle = \bigoplus_{y \in G} B_n(-gr(y))$. A *free B_n -persistence module* F is a B_n -persistence module such that for some set n -graded set G , $F \cong \langle G \rangle$.

Equivalently, we can define a free B_n -persistence module as a B_n -persistence module which satisfies a certain universal property. Free \mathbf{A}_n -persistence modules are defined via a universal property e.g. in [8, Section 4.2]. The definition for B_n -persistence modules is analogous; we refer the reader to [8] for details.

A *basis* for a free module F is a minimal set of generators for F . For G any graded set, identifying $y \in G$ with the copy of $1(-gr(y))$ in the summand $B_n(-gr(y))$ of $\langle G \rangle$ corresponding to y gives an identification of G with a basis for $\langle G \rangle$. It can be checked that if B and B' are two bases for a free B_n -persistence module F then $gr(B) = gr(B')$. Clearly then, $gr(B)$ of an arbitrarily chosen basis B for F is an isomorphism invariant of F and determines F up to isomorphism.

For R a homogeneous subset of a free B_n -persistence module F , $\langle R \rangle$ will always denote the submodule of F generated by R . Since, as noted above, R can be viewed as an n -graded set, we emphasize that for such R , $\langle R \rangle$ **does not** denote $\bigoplus_{y \in R} B_n(-gr(y))$.

2.5.3 Free Covers and Lifts

For M a B_n -persistence module, define a *free cover* of M be a pair (F_M, ρ_M) , where F_M is a free B_n -persistence module and $\rho_M : F_M \rightarrow M$ a surjective morphism of B_n -persistence modules.

For M, N B_n -persistence modules, (F_M, ρ_M) and (F_N, ρ_N) free covers of M and N , and $f : M \rightarrow N$ a morphism, define a *lift* of f to be a map $\tilde{f} : F_M \rightarrow F_N$ such that the following diagram commutes.

$$\begin{array}{ccc} F_M & \xrightarrow{\tilde{f}} & F_N \\ \downarrow \rho_M & & \downarrow \rho_N \\ M & \xrightarrow{f} & N \end{array}$$

Lemma 2.1 (Existence and Uniqueness up to Homotopy of Lifts). *For B_n -persistence modules M and N , free covers $(F_M, \rho_M), (F_N, \rho_N)$ of M, N , and a morphism $f : M \rightarrow N$, there exists a lift $\tilde{f} : F_M \rightarrow F_N$ of f . If $\tilde{f}' : F_M \rightarrow F_N$ is another lift of f , then $im(\tilde{f} - \tilde{f}') \subset ker(\rho_N)$.*

Proof. This is just a specialization of the standard result on the existence and homotopy uniqueness of free resolutions [25, Eisenbud A3.13] to the 0^{th} modules in free resolutions for M and N . The proof is straightforward. \square

2.5.4 Presentations of B_n -persistence Modules

A **presentation** of a B_n -persistence module M is a pair (G, R) where G is an n -graded set and $R \subset \langle G \rangle$ is a set of homogeneous elements such that $M \cong \langle G \rangle / \langle R \rangle$. We denote the presentation (G, R) as $\langle G | R \rangle$. For n -graded sets G_1, \dots, G_l and sets $R_1, \dots, R_m \subset \langle G_1 \cup \dots \cup G_l \rangle$, we'll let $\langle G_1, \dots, G_l | R_1, \dots, R_m \rangle$ denote $\langle G_1 \cup \dots \cup G_l | R_1 \cup \dots \cup R_m \rangle$.

If M is a B_n -persistence module such that there exists a presentation $\langle G | R \rangle$ for M with G and R finite, then we say M is *finitely presented*.

2.5.5 Minimal Presentations of B_n -persistence Modules

Let M be a B_n -persistence module. Define a presentation $\langle G | R \rangle$ of M to be *minimal* if

1. the quotient $\langle G \rangle \rightarrow \langle G \rangle / \langle R \rangle$ maps G to a minimal set of generators for $\langle G \rangle / \langle R \rangle$.
2. R is a minimal set of generators for $\langle R \rangle$.

It's clear that a minimal presentation for M exists.

Theorem 2.2. *If M is a finitely presented B_n -persistence module and $\langle G | R \rangle$ is a minimal presentation of M , then for any other presentation $\langle G' | R' \rangle$ of M , $gr(G) \leq gr(G')$ and $gr(R) \leq gr(R')$.*

Note that the theorem implies in particular that if $\langle G | R \rangle$ and $\langle G' | R' \rangle$ are two minimal presentations of M then $gr(G) = gr(G')$ and $gr(R) = gr(R')$.

We defer the proof of the theorem to Appendix B. The proof is an adaptation to our setting of a standard result [25, Theorem 20.2] about free resolutions of modules over a local ring. The main effort required in carrying out the adaptation is to prove that the ring B_n has a property known as *coherence*; we define coherence and prove that B_n is coherent in Appendix A.

3 Algebraic Preliminaries for 1-D Persistence

In this section, we review algebraic preliminaries and establish notation specific to 1-D persistent homology. This material will be used in Sections 4 and 5 to develop the machinery needed to prove Theorem 5.2.

3.0.6 Basic Notation

For S any subset of \mathbb{R}^2 , let $S_+ = \{(a, b) \in S | a < b\}$. For S a set and $f : S \rightarrow \mathbb{R}$ a function, let $\text{supp}(f) = \{s \in S | f(s) \neq 0\}$.

3.1 Structure Theorems For Tame \mathbf{A}_1 -Persistence Modules

The structure theorem for finitely generated \mathbf{A}_1 persistence modules [33] is well known in the applied topology community. In fact, this theorem generalizes to tame \mathbf{A}_1 -modules. The existence portion of the generalized theorem is given e.g. in [32]; the uniqueness is not mentioned there but is very easy to show; we do so below. To our knowledge, this generalization has

not previously been discussed in the computational topology literature. We will use the more general theorem to show that the bottleneck distance is equal to the interleaving distance for ordinary persistence.

Before stating the results, we establish some notation. For $a < b \in \mathbb{Z}$, Let $\mathbf{C}(a, b)$ denote the module $(k[x]/(x^{b-a}))(-a)$. Let $\mathbf{C}(a, \infty) = k[x](-a)$. Note that for fixed b (possibly infinite), the set of modules $\{\mathbf{C}(a, b)\}_{a \in (-\infty, b)}$ has a natural directed system structure; let $\mathbf{C}(-\infty, b)$ denote the colimit of this directed system.

For M a module and $m \in \mathbb{Z}_{\geq 0}$, let M^m denote the direct sum of m copies of M .

Theorem 3.1 (Structure Theorem for finitely generated \mathbf{A}_1 -persistence modules [33]). *Let \mathbf{M} be a finitely generated \mathbf{A}_1 -module. Then there is a unique function $\mathcal{D}_{\mathbf{M}} : (\mathbb{Z} \times \bar{\mathbb{Z}})_+ \rightarrow \mathbb{Z}_{\geq 0}$ with finite support such that*

$$\mathbf{M} \cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_{\mathbf{M}})} \mathbf{C}(a, b)^{\mathcal{D}_{\mathbf{M}}(a,b)}.$$

Theorem 3.2 (Structure Theorem for tame \mathbf{A}_1 -persistence modules [32]). *Let \mathbf{M} be a tame \mathbf{A}_1 -module. Then there is a unique function $\mathcal{D}_{\mathbf{M}} : \bar{\mathbb{Z}}_+^2 \rightarrow \mathbb{Z}_{\geq 0}$ such that*

$$\mathbf{M} \cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_{\mathbf{M}})} \mathbf{C}(a, b)^{\mathcal{D}_{\mathbf{M}}(a,b)}.$$

The uniqueness part of Theorem 3.2 is an immediate consequence of the following lemma, upon noting that the right hand sides of the equations in the statement of the lemma do not depend on $\mathcal{D}_{\mathbf{M}}$.

Lemma 3.3. *Let \mathbf{M} be a tame \mathbf{A}_1 -module, and let $\mathcal{D}_{\mathbf{M}} : \bar{\mathbb{Z}}_+^2 \rightarrow \mathbb{Z}_{\geq 0}$ be a function such that $\mathbf{M} \cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_{\mathbf{M}})} \mathbf{C}(a, b)^{\mathcal{D}_{\mathbf{M}}(a,b)}$. Then*

(i) *For $(a, b) \in \bar{\mathbb{Z}}_+^2$,*

$$\mathcal{D}_{\mathbf{M}}(a, b) = \text{rank}(\varphi_{\mathbf{M}}(a, b-1)) - \text{rank}(\varphi_{\mathbf{M}}(a, b)) - \text{rank}(\varphi_{\mathbf{M}}(a-1, b-1)) + \text{rank}(\varphi_{\mathbf{M}}(a-1, b)).$$

(ii) *For $b \in \mathbb{Z}$, $\mathcal{D}_{\mathbf{M}}(-\infty, b) = \lim_{a \rightarrow -\infty} \text{rank}(\varphi_{\mathbf{M}}(a, b-1)) - \lim_{a \rightarrow -\infty} \text{rank}(\varphi_{\mathbf{M}}(a, b))$.*

(iii) *For $a \in \mathbb{Z}$, $\mathcal{D}_{\mathbf{M}}(a, \infty) = \lim_{b \rightarrow \infty} \text{rank}(\varphi_{\mathbf{M}}(a, b)) - \lim_{b \rightarrow \infty} \text{rank}(\varphi_{\mathbf{M}}(a-1, b))$.*

(iv) *$\mathcal{D}_{\mathbf{M}}(-\infty, \infty) = \lim_{a \rightarrow -\infty} \lim_{b \rightarrow \infty} \text{rank}(\varphi_{\mathbf{M}}(a, b))$.*

Proof. This is trivial. □

We call $\mathcal{D}_{\mathbf{M}}$ the **(discrete) persistence diagram** of \mathbf{M} .

In Section 4, we prove a structure theorem analogous to Theorem 3.2 for a subset of the tame B_1 -persistence modules which contains the finitely presented B_1 -persistence modules. We do not address the problem of generalizing this structure theorem to the full set of tame B_1 -persistence modules, but to echo a sentiment expressed in [11], it would be nice to have such a result.

3.2 Discrete Persistence modules.

In order to define persistence diagrams of B_1 -persistence modules, we need a mild generalization of \mathbf{A}_1 -persistence modules.

Let $S \subset \mathbb{R}$ be a countably infinite set with no accumulation point. The authors of [11] define a **discrete persistence module** M_S to be a collection of vector spaces $\{M_s\}_{s \in S}$ indexed by S together with linear maps $\{\varphi_{M_S}(s_1, s_2)\}_{s_1 \leq s_2 \in S}$.

Define a **grid function** $t : \mathbb{Z} \rightarrow \mathbb{R}$ to be a strictly increasing function with no accumulation point.

Remark 3.1. Discrete persistence modules are of course closely related to \mathbf{A}_1 -persistence modules. A countably infinite subset of $S \subset \mathbb{R}$ with no accumulation point can be indexed by a grid function t with image S , and such a grid function is uniquely determined by the value of $t(0)$. Thus, pairs (M_S, s) , where M_S is a discrete persistence module and s is an element of S , are equivalent to pairs (\mathbf{M}', t) , where \mathbf{M}' is an \mathbf{A}_1 -persistence module and t is a grid function; there is an equivalence sending each pair (M_S, s) to the pair (\mathbf{M}', t) , where t is a grid function with $\text{im}(t) = S$, $t(0) = s$, and \mathbf{M}' is the \mathbf{A}_1 -persistence module such that for $z \in \mathbb{Z}$, $\mathbf{M}'_z = M_{t(z)}$ and $\varphi_{\mathbf{M}'}(z_1, z_2) = \varphi_{M_S}(t(z_1), t(z_2))$.

As a matter of expository convenience, from now on we'll define discrete modules to be pairs (\mathbf{M}, t) where \mathbf{M} is an \mathbf{A}_1 -persistence module and t is a grid function. This in effect means we are carrying around the extra data of an element of S in our discrete persistent modules relative to those defined in [11], but this won't present a problem—in particular, the definition of the persistence diagram of a discrete persistence module that we present below is independent of the choice of this element, and is equivalent to that of [11].

3.3 Persistence Diagrams

The definition of a persistence diagram that we present here differs in some cosmetic respects from that in [11]. Our choice in this regard is a matter of notational convenience; the reader may check that our definition of the bottleneck distance between tame B_1 -persistence modules is equivalent to that of [11].

For a grid function t , define $\bar{t} : \bar{\mathbb{Z}} \rightarrow \bar{\mathbb{R}}$ as

$$\bar{t}(z) = \begin{cases} t(z) & \text{if } z \in \mathbb{Z}, \\ -\infty & \text{if } z = -\infty, \\ \infty & \text{if } z = \infty. \end{cases}$$

Let $\bar{t} \times \bar{t} : \bar{\mathbb{Z}}_+^2 \rightarrow \bar{\mathbb{R}}_+^2$ be defined by $\bar{t} \times \bar{t}(a, b) = (\bar{t}(a), \bar{t}(b))$.

We define a **persistence diagram** to be a function $\mathcal{D} : \bar{\mathbb{R}}_+^2 \rightarrow \mathbb{Z}_{\geq 0}$.

For (\mathbf{M}, t) a discrete persistence module, define $\mathcal{D}_{(\mathbf{M}, t)}$, the **persistence diagram of (\mathbf{M}, t)** , to be the persistence diagram for which $\text{supp}(\mathcal{D}_{(\mathbf{M}, t)}) = \bar{t} \times \bar{t}(\text{supp}(\mathcal{D}_{\mathbf{M}}))$ and so that $\mathcal{D}_{(\mathbf{M}, t)}(\bar{t}(a), \bar{t}(b)) = \mathcal{D}_{\mathbf{M}}(a, b)$ for all $(a, b) \in \bar{\mathbb{Z}}_+^2$.

3.3.1 Bottleneck Metric

For $x \in \mathbb{R}$, define $x + \infty = \infty$ and $x - \infty = -\infty$. Then the usual definition of l^∞ norm on the plane extends to $\bar{\mathbb{R}}^2$; we denote it by $\|\cdot\|_\infty$.

Now define a **multibijection** between two persistence diagrams $\mathcal{D}_1, \mathcal{D}_2$ to be a function $\gamma : \text{supp}(\mathcal{D}_1) \times \text{supp}(\mathcal{D}_2) \rightarrow \mathbb{Z}_{\geq 0}$ such that

1. For each $x \in \text{supp}(\mathcal{D}_1)$, the set $\{y \in \text{supp}(\mathcal{D}_2) | (x, y) \in \text{supp}(\gamma)\}$ is finite and

$$\mathcal{D}_1(x) = \sum_{y \in \text{supp}(\mathcal{D}_2)} \gamma(x, y),$$

2. For each $y \in \text{supp}(\mathcal{D}_2)$, the set $\{x \in \text{supp}(\mathcal{D}_1) | (x, y) \in \text{supp}(\gamma)\}$ is finite and

$$\mathcal{D}_2(y) = \sum_{x \in \text{supp}(\mathcal{D}_1)} \gamma(x, y).$$

For persistence diagrams $\mathcal{D}_1, \mathcal{D}_2$, let $\mathcal{L}(\mathcal{D}_1, \mathcal{D}_2)$ denote the triples $(\mathcal{D}'_1, \mathcal{D}'_2, \gamma)$, where \mathcal{D}'_1 and \mathcal{D}'_2 are persistence diagrams with $\mathcal{D}'_1 \leq \mathcal{D}_1, \mathcal{D}'_2 \leq \mathcal{D}_2$, and γ is a multibijection between \mathcal{D}'_1 and \mathcal{D}'_2 .

We define the **bottleneck metric** d_B between two persistence diagrams $\mathcal{D}_1, \mathcal{D}_2$ as

$$d_B(\mathcal{D}_1, \mathcal{D}_2) = \inf_{\substack{(\mathcal{D}'_1, \mathcal{D}'_2, \gamma) \\ \in \mathcal{L}(\mathcal{D}_1, \mathcal{D}_2)}} \max \left(\sup_{\substack{(a,b) \in \text{supp}(\mathcal{D}_1 - \mathcal{D}'_1) \\ \cup \text{supp}(\mathcal{D}_2 - \mathcal{D}'_2)}} \frac{1}{2}(b - a), \sup_{(x,y) \in \text{supp}(\gamma)} \|y - x\|_\infty \right).$$

3.3.2 Discretizations of B_1 -modules

Let t be a grid function. For M a B_1 -persistence module, we define the t -discretization of M to be the discrete persistence module $(\mathbf{P}_t(M), t)$ with $\mathbf{P}_t(M)$ defined as follows:

1. For $z \in \mathbb{Z}$, $\mathbf{P}_t(M)_z = M_{t(z)}$; let $\mathcal{I}_{M,t,z} : \mathbf{P}_t(M)_z \rightarrow M_{t(z)}$ denote this identification.
2. For $y, z \in \mathbb{Z}$, $y \leq z$, $\varphi_{\mathbf{P}_t(M)}(y, z) = \mathcal{I}_{M,t,z}^{-1} \circ \varphi_M(t(y), t(z)) \circ \mathcal{I}_{M,t,y}$.

3.3.3 Persistence diagrams of B_1 -persistence modules

We'll say a grid function t is an ϵ -**cover** if for any $a \in \mathbb{R}$, there exists $b \in \text{im}(t)$ with $|a - b| \leq \epsilon$. Now fix $\alpha \in \mathbb{R}$ and let $\{t_i\}_{i=1}^\infty$ be a sequence of grid functions with t_i a $1/2^i$ -cover.

It is asserted in [11] that for any tame B_1 -persistence module M the persistence diagrams $\mathcal{D}_{(\mathbf{P}_{t_i}(M), t_i)}$ converge in the bottleneck metric to a limiting persistence diagram \mathcal{D}_M and that \mathcal{D}_M is independent of the choice of the sequence $\{t_i\}$. We call \mathcal{D}_M the persistence diagram of M . For M and N tame B_1 -persistence modules, we define $d_B(M, N) = d_B(\mathcal{D}_M, \mathcal{D}_N)$.

Remark 3.2. Two non-isomorphic tame B_1 -persistence modules can have identical persistence diagrams. For example, take M and N to be the B_1 -persistence modules of Example 2.1. M and N are not isomorphic but it is easy to check that they have the same persistence diagram. Thus d_B defines a pseudometric (but not a metric) on isomorphism classes of tame B_1 -persistence modules.

4 Structure Theorem for Well Behaved B_1 -persistence modules

In this section we prove an analogue of Theorem 3.2 for a certain subset of the tame B_1 -persistence modules which we call the *well behaved* persistence modules. The set of well behaved persistence modules contains the set of finitely presented B_1 -persistence modules. These modules are in a sense “essentially discrete.” Indeed, they are exactly the B_1 -persistence modules that are the images of tame \mathbf{A}_1 -persistence modules under a certain family of functors from $\mathbf{A}_1\text{-mod}$ to $B_1\text{-mod}$.

Our strategy for proving the structure theorem for well behaved persistence modules is to exploit Theorem 3.2, taking advantage of the functorial relationship between \mathbf{A}_1 -persistence modules and well behaved B_1 -persistence modules.

4.1 Well Behaved Persistence Modules

A *critical value* of a B_1 -persistence module M is a point $a \in \mathbb{R}$ such that for no $\epsilon \in \mathbb{R}_{\geq 0}$ is it true that for all $u \leq v \in [a - \epsilon, a + \epsilon]$, $\varphi_M(u, v)$ is an isomorphism.

We’ll say a tame B_1 -persistence module M is **well behaved** if

1. The critical values of M are countable and have no accumulation point.
2. For each critical point a of M , there exists $\epsilon > 0$ such that $\varphi_M(a, y)$ is an isomorphism for all $y \in [a, a + \epsilon]$.

Proposition 4.1. *A finitely presented B_1 -persistence module is well behaved.*

Proof. Let M be a finitely presented B_1 -persistence module and let $U \subset \mathbb{R}$ be the set of grades of the generators and relations in a minimal presentation for M . (It follows from Theorem 2.2 that U is well defined). Lemma 6.4 below tells us that for any $a \leq b \in \mathbb{R}$ such that $(a, b] \cap U = \emptyset$, $\varphi_M(a, b)$ is an isomorphism. Since U is finite, the result follows immediately. \square

Let t be a grid function. Define $t^{-1} : \mathbb{R} \rightarrow \mathbb{Z}$ by $t^{-1}(y) = \max\{z \in \mathbb{Z} | t(z) \leq y\}$.

Define $\bar{t}^{-1} : \bar{\mathbb{R}} \rightarrow \bar{\mathbb{Z}}$ by

$$\bar{t}^{-1}(u) = \begin{cases} t^{-1}(u) & \text{if } u \in \mathbb{R}, \\ -\infty & \text{if } u = -\infty, \\ \infty & \text{if } u = \infty. \end{cases}$$

We’ll now define a functor $E_t : \mathbf{A}_1\text{-mod} \rightarrow B_1\text{-mod}$ as follows:

1. Action of E_t on objects: For \mathbf{M} an \mathbf{A}_1 -persistence module and $u \in \mathbb{R}$, $E_t(\mathbf{M})_u = \mathbf{M}_{t^{-1}(u)}$; let $\mathcal{J}_{\mathbf{M}, t, u} : E_t(\mathbf{M})_u \rightarrow \mathbf{M}_{t^{-1}(u)}$ denote this identification. For $u, v \in \mathbb{R}$, $u \leq v$, let $\varphi_{E_t(\mathbf{M})}(u, v) = \mathcal{J}_{\mathbf{M}, t, v}^{-1} \circ \varphi_{\mathbf{M}}(t^{-1}(u), t^{-1}(v)) \circ \mathcal{J}_{\mathbf{M}, t, u}$.
2. Action of E_t on morphisms: For \mathbf{M} and \mathbf{N} \mathbf{A}_1 -persistence modules and $f \in \text{hom}(\mathbf{M}, \mathbf{N})$, define $E_t(f) : E_t(\mathbf{M}) \rightarrow E_t(\mathbf{N})$ by letting $E_t(f)_u = \mathcal{J}_{\mathbf{N}, t, u}^{-1} \circ f_{t^{-1}(u)} \circ \mathcal{J}_{\mathbf{M}, t, u}$ for all $u \in \mathbb{R}$.

We leave to the reader the easy verification that E_t is in fact a functor with target $B_1\text{-mod}$.

It's clear that if \mathbf{M} is a tame \mathbf{A}_1 -persistence module, then for any grid function t , $E_t(\mathbf{M})$ is tame. Moreover, it's easy to check that for any grid function t and any tame \mathbf{A}_1 -persistence module \mathbf{M} , $E_t(\mathbf{M})$ is well behaved.

Conversely, we have the following:

Proposition 4.2. *If M is a well behaved B_1 -persistence module, then there is some tame \mathbf{A}_1 -persistence-module \mathbf{M} and some grid function t such that $M \cong E_t(\mathbf{M})$.*

Proof. Let $t : \mathbb{Z} \rightarrow \mathbb{R}$ be a grid function whose image contains the critical points of M . Let (\mathbf{M}, t) denote the t -discretization of M , as defined in Section 3.3.2. \mathbf{M} clearly is tame. We'll show that $M \cong E_t(\mathbf{M})$.

For $u \in \mathbb{R}$, define $\sigma_u : E_t(\mathbf{M})_u \rightarrow M_u$ by $\sigma_u = \varphi_M(t \circ t^{-1}(u), u) \circ \mathcal{I}_{M,t,t^{-1}(u)} \circ \mathcal{J}_{\mathbf{M},t,u}$. By definition, $\mathcal{J}_{\mathbf{M},t,u}$ and $\mathcal{I}_{M,t,t^{-1}(u)}$ are isomorphisms. Moreover, a simple compactness argument shows that since M is well behaved, $\varphi_M(t \circ t^{-1}(u), u)$ is an isomorphism. Thus σ_u is an isomorphism.

We claim that the collection of maps $\{\sigma_u\}_{u \in \mathbb{R}}$ defines an isomorphism of modules. To see this, we need to show that for all $u, v \in \mathbb{R}, u \leq v$, $\sigma_v \circ \varphi_{E_t(\mathbf{M})}(u, v) = \varphi_M(u, v) \circ \sigma_u$.

$$\begin{aligned}
\sigma_v \circ \varphi_{E_t(\mathbf{M})}(u, v) &= \sigma_v \circ \mathcal{J}_{\mathbf{M},t,v}^{-1} \circ \varphi_{\mathbf{M}}(t^{-1}(u), t^{-1}(v)) \circ \mathcal{J}_{\mathbf{M},t,u} \\
&= \varphi_M(t \circ t^{-1}(v), v) \circ \mathcal{I}_{M,t,t^{-1}(v)} \circ \mathcal{J}_{\mathbf{M},t,v} \circ \mathcal{J}_{\mathbf{M},t,v}^{-1} \circ \varphi_{\mathbf{M}}(t^{-1}(u), t^{-1}(v)) \circ \mathcal{J}_{\mathbf{M},t,u} \\
&= \varphi_M(t \circ t^{-1}(v), v) \circ \mathcal{I}_{M,t,t^{-1}(v)} \circ \varphi_{\mathbf{M}}(t^{-1}(u), t^{-1}(v)) \circ \mathcal{J}_{\mathbf{M},t,u} \\
&= \varphi_M(t \circ t^{-1}(v), v) \circ \mathcal{I}_{M,t,t^{-1}(v)} \circ \mathcal{I}_{M,t,t^{-1}(v)}^{-1} \\
&\quad \circ \varphi_M(t \circ t^{-1}(u), t \circ t^{-1}(v)) \circ \mathcal{I}_{M,t,t^{-1}(u)} \circ \mathcal{J}_{\mathbf{M},t,u} \\
&= \varphi_M(t \circ t^{-1}(v), v) \circ \varphi_M(t \circ t^{-1}(u), t \circ t^{-1}(v)) \circ \mathcal{I}_{M,t,t^{-1}(u)} \circ \mathcal{J}_{\mathbf{M},t,u} \\
&= \varphi_M(t \circ t^{-1}(u), v) \circ \mathcal{I}_{M,t,t^{-1}(u)} \circ \mathcal{J}_{\mathbf{M},t,u} \\
&= \varphi_M(u, v) \circ \varphi_M(t \circ t^{-1}(u), u) \circ \mathcal{I}_{M,t,t^{-1}(u)} \circ \mathcal{J}_{\mathbf{M},t,u} \\
&= \varphi_M(u, v) \circ \sigma_u.
\end{aligned}$$

□

Remark 4.1. The material above can be adapted with only minor changes to the setting of B_n -persistence modules, where it sheds some light on the relationship between \mathbf{A}_n -persistence modules and B_n -persistence modules. Namely, the definitions of a well behaved persistence module, grid function, and the functors E_t generalize to the multidimensional setting, and analogues of Propositions 4.1 and 4.2 hold in that setting. It can be shown that the functor $(\cdot) \otimes_{\mathbf{A}_n} B_n$ mentioned in Section 2.3.3 is naturally isomorphic to a generalized functor E_t . The generalization of the above material also can be used to translate algebraic results about \mathbf{A}_n -persistence modules into analogous results about B_n -persistence modules. For example, it can be used to show that any finitely presented B_n -persistence module has a free resolution of

length at most n —that is, an analogue of the Hilbert syzygy theorem holds for B_n -persistence modules.

However, as we have no immediate need for the generalization or its consequences in this paper, we omit it.

4.2 The Structure Theorem

First, note that for any $a \in \mathbb{R}_{\geq 0}$, $k[[a, \infty]]$ (as defined in Section 2.2) is an ideal of B_1 .

For $a < b \in \mathbb{R}$, let $C(a, b)$ denote $(B_1/k[[b - a, \infty]])(-a)$; let $C(a, \infty)$ denote $B_1(-a)$. In analogy to the discrete case, for fixed b (possibly infinite), the set of modules $\{C(a, b) | a \in \mathbb{R}, a < b\}$ has a natural directed system structure; let $C(-\infty, b)$ denote the colimit of this directed system.

Lemma 4.3. *Let M be a well-behaved persistence module and let \mathcal{D} be a persistence diagram such that $M \cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D})} C(a, b)^{\mathcal{D}(a,b)}$. Then $\mathcal{D}_M = \mathcal{D}$.*

Proof. Let

$$A = \{a \in \mathbb{R} | (a, b) \in \text{supp}(\mathcal{D}) \text{ for some } b \in \bar{\mathbb{R}}\} \cup \{b \in \mathbb{R} | (a, b) \in \text{supp}(\mathcal{D}) \text{ for some } a \in \bar{\mathbb{R}}\}.$$

Let t be a grid function such that $A \subset \text{im}(t)$. We claim that $\mathcal{D}_{(\mathbf{P}_t(M), t)} = \mathcal{D}$. Since $\text{supp}(\mathcal{D}) \in \text{im}(\bar{t} \times \bar{t})$, this is true if and only if $\mathcal{D}_{\mathbf{P}_t(M)}(y, z) = \mathcal{D}(\bar{t}(y), \bar{t}(z))$ for all $(y, z) \in \bar{\mathbb{Z}}_+^2$.

To show that $\mathcal{D}_{\mathbf{P}_t(M)}(y, z) = \mathcal{D}(\bar{t}(y), \bar{t}(z))$ for all $(y, z) \in \bar{\mathbb{Z}}_+^2$, we'll need the following analogue of Lemma 3.3.

Lemma 4.4. *Let M, \mathcal{D} , and t be as above.*

(i) For $(y, z) \in \mathbb{Z}_+^2$,

$$\begin{aligned} \mathcal{D}(t(y), t(z)) &= \text{rank}(\varphi_M(t(y), t(z-1))) - \text{rank}(\varphi_M(t(y), t(z))) \\ &\quad - \text{rank}(\varphi_M(t(y-1), t(z-1))) + \text{rank}(\varphi_M(t(y-1), t(z))). \end{aligned}$$

(ii) For $z \in \mathbb{Z}$,

$$\mathcal{D}(-\infty, t(z)) = \lim_{y \rightarrow -\infty} \text{rank}(\varphi_M(t(y), t(z-1))) - \lim_{y \rightarrow -\infty} \text{rank}(\varphi_M(t(y), t(z))).$$

(iii) For $y \in \mathbb{Z}$,

$$\mathcal{D}(t(y), \infty) = \lim_{z \rightarrow \infty} \text{rank}(\varphi_M(t(y), t(z))) - \lim_{b \rightarrow \infty} \text{rank}(\varphi_M(t(y-1), t(z))).$$

(iv)

$$\mathcal{D}(-\infty, \infty) = \lim_{y \rightarrow -\infty} \lim_{z \rightarrow \infty} \text{rank}(\varphi_M(t(y), t(z))).$$

Proof. The proof is straightforward; we omit it. □

For $(y, z) \in \mathbb{Z}_+^2$ we have

$$\begin{aligned}
\mathcal{D}_{\mathbf{P}_t(M)}(y, z) &= \text{rank}(\varphi_{\mathbf{P}_t(M)}(y, z-1)) - \text{rank}(\varphi_{\mathbf{P}_t(M)}(y, z)) \\
&\quad - \text{rank}(\varphi_{\mathbf{P}_t(M)}(y-1, z-1)) + \text{rank}(\varphi_{\mathbf{P}_t(M)}(y-1, z)) \\
&= \text{rank}(\varphi_M(t(y), t(z-1))) - \text{rank}(\varphi_M(t(y), t(z))) \\
&\quad - \text{rank}(\varphi_M(t(y-1), t(z-1))) + \text{rank}(\varphi_M(t(y-1), t(z))) \\
&= \mathcal{D}(t(y), t(z)),
\end{aligned}$$

where the first equality follows from Lemma 3.3(i), and the last equality follows from Lemma 4.4(i).

Thus we have $\mathcal{D}_{\mathbf{P}_t(M)}(y, z) = \mathcal{D}(\bar{t}(y), \bar{t}(z))$ for all $(y, z) \in \mathbb{Z}_+^2$. Similar arguments using Lemma 3.3(ii)-(iv) and Lemma 4.4(ii)-(iv) in the cases where $y = -\infty$ or $z = \infty$ show that in fact this holds for $(y, z) \in \mathbb{Z}_+^2$. This proves the claim.

It follows easily from the fact that M is well behaved that A is equal to the set of critical values of M . There thus exists a sequence of grid functions $\{t_i\}_{i \in \mathbb{N}}$ such that t_i is a $1/2^i$ cover and $A \subset \text{im}(t_i)$ for each i . The lemma follows by writing \mathcal{D}_M as the limit of the persistence diagrams $\mathcal{D}_{(\mathbf{P}_{t_i}(M), t_i)}$. \square

Theorem 4.5. *Let M be a well behaved B_1 -persistence module. Let \mathcal{D}_M be the persistence diagram of M . Then*

$$M \cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_M)} C(a, b)^{\mathcal{D}_M(a,b)}.$$

This decomposition of M is unique in the sense that if \mathcal{D} is another persistence diagram such that $M \cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D})} C(a, b)^{\mathcal{D}(a,b)}$, then $\mathcal{D} = \mathcal{D}_M$.

Proof. By Lemma 4.3, it's enough show that there exists some persistence diagram \mathcal{D} such that $M \cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D})} C(a, b)^{\mathcal{D}(a,b)}$.

By Proposition 4.2, there exists a grid function t and a tame \mathbf{A}_1 -persistence module \mathbf{M} such that $E_t(\mathbf{M}) \cong M$. The structure theorem for tame \mathbf{A}_1 -persistence modules gives us that there's a persistence diagram $\mathcal{D}_{\mathbf{M}}$ supported in $\bar{\mathbb{Z}}_+^2$ such that we may take $\mathbf{M} = \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_{\mathbf{M}})} \mathbf{C}(a, b)^{\mathcal{D}_{\mathbf{M}}(a,b)}$. We'll show that $M \cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_{\mathbf{M}})} E_t(\mathbf{C}(a, b))^{\mathcal{D}_{\mathbf{M}}(a,b)}$. We have that $E_t(\mathbf{C}(a, b)) \cong C(\bar{t}(a), \bar{t}(b))$ for any $(a, b) \in \bar{\mathbb{Z}}_+^2$, so this gives the result.

To show that $M \cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_{\mathbf{M}})} E_t(\mathbf{C}(a, b))^{\mathcal{D}_{\mathbf{M}}(a,b)}$, we'll use the category theoretic characterization of direct sums of modules as coproducts [28]. Recall that in an arbitrary category, an object X is a *coproduct* of objects $\{X^\alpha\}_{\alpha \in A}$ iff there exist morphisms $\{i^\alpha : X^\alpha \rightarrow X\}_{\alpha \in A}$, called canonical injections, with the following universal property: for any object Y and morphisms $\{f^\alpha : X^\alpha \rightarrow Y\}_{\alpha \in A}$, there exists a unique morphism $f : X \rightarrow Y$ such that $f \circ i^\alpha = f^\alpha$ for each $\alpha \in A$. In a category of modules over a ring R , The coproduct of modules X^α is $\bigoplus_\alpha X^\alpha$; the canonical injections are just the usual inclusions $X^\alpha \hookrightarrow \bigoplus_\alpha X^\alpha$. The same is thus true for the module subcategories $\mathbf{A}_n\text{-mod}$ and $B_n\text{-mod}$.

Now let $\{\mathbf{M}^\alpha\}$ denote the indecomposable summands of \mathbf{M} in the direct sum decomposition $\mathbf{M} = \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_{\mathbf{M}})} \mathbf{C}(a, b)^{\mathcal{D}_{\mathbf{M}}(a,b)}$, so that each $\mathbf{M}^\alpha = \mathbf{C}(a, b)$ for some $(a, b) \in \bar{\mathbb{Z}}_+^2$. Let $\{i^\alpha : \mathbf{M}^\alpha \rightarrow \mathbf{M}\}$ denote the canonical injections.

We'll show that the maps $E_t(i^\alpha) : E_t(\mathbf{M}^\alpha) \rightarrow E_t(\mathbf{M})$ satisfy the universal property of a coproduct, so that $M \cong E_t(\mathbf{M}) \cong \oplus_\alpha E_t(\mathbf{M}^\alpha)$ as desired.

To show that the maps $E_t(i^\alpha) : E_t(\mathbf{M}^\alpha) \rightarrow E_t(\mathbf{M})$ satisfy the universal property of a coproduct, let Y be an arbitrary B_1 -persistence module and $\{f^\alpha : E_t(\mathbf{M}^\alpha) \rightarrow Y\}$ be homomorphisms.

For any $z \in \mathbb{Z}$, $\mathbf{M}_z \cong \oplus_\alpha \mathbf{M}_z^\alpha$. It follows from the definition of E_t that for any $r \in \mathbb{R}$,

$$E_t(\mathbf{M})_r \cong \oplus_\alpha E_t(\mathbf{M}^\alpha)_r$$

with the maps $E_t(i^\alpha)_r$ the canonical inclusions.

For each $r \in \mathbb{R}$, define $f_r : E_t(\mathbf{M})_r \rightarrow Y_r$ as $\oplus_\alpha f_r^\alpha$ (i.e. f_r is the map guaranteed to exist by the universal property of direct sums for vector spaces.) It's easy to check that the maps f_r commute with the transition maps in $E_t(\mathbf{M})$ and Y , so that they define a morphism $f : E_t(\mathbf{M}) \rightarrow Y$. We also have that $f \circ E_t(i^\alpha) = f^\alpha$ for each α . By the universal property of direct sums of vector spaces, for each r f_r is the unique linear transformation from $E_t(\mathbf{M})_r$ to Y_r such that for each α , $f_r \circ E_t(i^\alpha)_r = f_r^\alpha$. Therefore f must itself satisfy the desired uniqueness property. This completes the proof. \square

5 The Equality of d_I and d_B on Tame B_1 -persistence Modules

We show in this section that the restriction of the interleaving distance to tame B_1 -persistence modules is equal to the bottleneck distance. This shows that the interleaving distance is in fact a generalization of the bottleneck distance, as we want. The result is also instrumental in proving Corollary 6.3, our converse to the algebraic stability result of [11].

5.0.1 The Algebraic Stability of Persistence

The main result of [11], generalizing considerably the earlier result of [17], is the following:

Theorem 5.1 (Algebraic Stability of Persistence). *Let $\epsilon > 0$, and let M and N be two tame B_1 -persistence modules. If M and N are ϵ -interleaved, then $d_B(M, N) \leq \epsilon$.*

5.0.2 A Converse to the Algebraic Stability of Persistence?

In the conclusion of [11], the authors ask whether it's true that if M and N are tame B_1 -persistence modules with $d_B(M, N) = \epsilon$ then M and N are ϵ -interleaved. Example 2.1 shows that this is not true. However, Corollary 6.3 below, which follows immediately from Theorems 5.2 and 6.1, asserts that the result is true provided M and N are finitely presented. More generally, Corollary 6.3 tells us that if M and N are tame modules with $d_B(M, N) = \epsilon$, then M and N are $(\epsilon + \delta)$ -interleaved for any $\delta > 0$. In other words, the converse of Theorem 5.1 holds for tame modules to arbitrarily small error.

Theorem 5.2. $d_B(M, N) = d_I(M, N)$ for any tame B_1 -persistence modules M and N .

Proof. Theorem 5.1 tells us that $d_B(M, N) \leq d_I(M, N)$, so we just need to show that $d_B(M, N) \geq d_I(M, N)$. It will follow from the structure theorem for well behaved persistence modules (Theorem 4.5) that $d_B(M', N') \geq d_I(M', N')$ for well behaved persistence modules M' and N' (Lemma 5.4 below). To extend this result to arbitrary tame modules, we will approximate the modules M and N up to arbitrarily small error in the interleaving distance by well behaved persistence modules (Lemma 5.5 below). The full result will follow readily from this approximation.

Lemma 5.3. *If $(a, b), (a', b') \in \bar{\mathbb{R}}_+^2$ with $\|(a, b) - (a', b')\|_\infty \leq \epsilon$, then $C(a, b)$ and $C(a', b')$ are ϵ -interleaved.*

Proof. This is easy to prove; we leave the details to the reader. \square

Lemma 5.4. *If M and N are two well behaved persistence modules and $d_B(M, N) = \epsilon$ then $d_I(M, N) = \epsilon$.*

Proof. By stability we just need to show that $d_I(M, N) \leq \epsilon$. By the structure theorem for well behaved persistence modules (Theorem 3.2), we have that

$$\begin{aligned} M &\cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_M)} C(a, b)^{\mathcal{D}_M(a,b)}, \\ N &\cong \bigoplus_{(a,b) \in \text{supp}(\mathcal{D}_N)} C(a, b)^{\mathcal{D}_N(a,b)}. \end{aligned}$$

Since $d_B(M', N') = \epsilon$, for any $\delta > 0$ there exist persistence diagrams D'_M and D'_N with $D'_M \leq D_M$, $D'_N \leq D_N$, and a multibijection γ between D'_M and D'_N such that

1. For any $(a, b) \in \text{supp}(D_M - D'_M) \cup \text{supp}(D_N - D'_N)$, $(b - a)/2 \leq \epsilon + \delta$,
2. For any $(x, y) \in \text{supp}(\gamma)$, $\|x - y\|_\infty \leq \epsilon + \delta$.

Fix such D'_M , D'_N , and γ . Now we can choose well behaved modules $M', M'' \subset M$ and $N', N'' \subset N$ such that $M = M' \oplus M''$, $N = N' \oplus N''$, $D_{M'} = D'_M$, $D_{N'} = D'_N$, $D_{M''} = D_M - D'_M$, and $D_{N''} = D_N - D'_N$.

It follows from Lemma 5.3 that for each $(a, b), (a', b') \in \text{supp}(\gamma)$, $C(a, b)$ and $C(a', b')$ are $(\epsilon + \delta)$ -interleaved. We may write

$$\begin{aligned} M' &\cong \bigoplus_{(a,b),(a',b') \in \text{supp}(\gamma)} C(a, b)^{\gamma((a,b),(a',b'))}, \\ N' &\cong \bigoplus_{(a,b),(a',b') \in \text{supp}(\gamma)} C(a', b')^{\gamma((a,b),(a',b'))}. \end{aligned}$$

It's clear from the form of these decompositions for M' and N' that a choice of a pair of $(\epsilon + \delta)$ -interleaving homomorphisms between $C(a, b)$ and $C(a', b')$ for each pair $(a, b), (a', b') \in \text{supp}(\gamma)$ induces a pair of $(\epsilon + \delta)$ -interleaving homomorphisms $\hat{f} : M' \rightarrow N'(\epsilon + \delta)$ and $\hat{g} : N' \rightarrow M'(\epsilon + \delta)$.

Now we extend this pair to a pair of homomorphisms $f : M \rightarrow N(\epsilon + \delta)$, $g : N \rightarrow M(\epsilon + \delta)$ by defining $f(y) = \hat{f}(y)$ for $y \in M'$, $f(M'') = 0$, $g(y) = \hat{g}(y)$ for $y \in N'$, and $g(N'') = 0$. Obviously, f and g restrict to $(\epsilon + \delta)$ -interleaving homomorphisms between M' and N' . Moreover, we have that $S(M'', 2\epsilon + \delta) = 0$ and $S(N'', 2\epsilon + \delta) = 0$, so f and g restrict to $(\epsilon + \delta)$ -interleaving homomorphisms between M'' and N'' . Thus by linearity, f and g are $(\epsilon + \delta)$ -interleaving homomorphisms between M and N . Since δ may be taken to be arbitrarily small, we must have $d_I(M, N) \leq \epsilon$, as we wanted to show. \square

Lemma 5.5. *For any tame B_1 -persistence module M and $\delta > 0$, there exists a well behaved persistence module M' with $d_I(M, M') \leq \delta$.*

Proof. Let t be an $\delta/2$ -cover of \mathbb{R} , as defined in Section 3.3.3. For any $r \in \mathbb{R}$, there exists $r' \in \text{im}(t)$, with $0 \leq r' - r \leq \delta$. Define a function $\lambda : \mathbb{R} \rightarrow \text{im}(t)$ such that $\lambda(r) = \min\{r' \geq r \mid r' \in \text{im}(t)\}$. Then $0 \leq \lambda(r) - r \leq \delta$ for all $r \in \mathbb{R}$.

Let $\mathbf{M} = \mathbf{P}_t(M)$ (as defined in Section 3.3.2) and let $M' = E_t(\mathbf{M})$. Then M' is well-behaved. We now show that M and M' are δ -interleaved, which implies that $d_I(M, M') \leq \delta$.

Define $f : M \rightarrow M'(\delta)$ to be the morphism for which

$$f_u : M_u \rightarrow M'_{u+\delta} = \varphi_{M'}(\lambda(u), u + \delta) \circ \mathcal{J}_{\mathbf{M}, t, \lambda(u)}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(\lambda(u))}^{-1} \circ \varphi_M(u, \lambda(u)).$$

Define $g : M' \rightarrow M(\delta)$ to be the morphism for which

$$g_u : M'_u \rightarrow M_{u+\delta} = \varphi_M(\lambda(u), u + \delta) \circ \mathcal{I}_{M, t, t^{-1}(\lambda(u))} \circ \mathcal{J}_{\mathbf{M}, t, \lambda(u)} \circ \varphi_{M'}(u, \lambda(u)).$$

We need to check that f and g thus defined are in fact morphisms. We verify this for f ; the verification for g is similar; we omit it.

If $u \leq v \in \mathbb{R}$, we have

$$f_v \circ \varphi_M(u, v) = \varphi_{M'}(\lambda(v), v + \delta) \circ \mathcal{J}_{\mathbf{M}, t, \lambda(v)}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(\lambda(v))}^{-1} \circ \varphi_M(v, \lambda(v)) \circ \varphi_M(u, v). \quad (1)$$

By definition, for any $u \leq v \in \mathbb{R}$, we have

$$\begin{aligned} \varphi_{M'}(u, v) &= \mathcal{J}_{\mathbf{M}, t, v}^{-1} \circ \varphi_{\mathbf{M}}(t^{-1}(u), t^{-1}(v)) \circ \mathcal{J}_{\mathbf{M}, t, u} \\ &= \mathcal{J}_{\mathbf{M}, t, v}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(v)}^{-1} \circ \varphi_M(u, v) \circ \mathcal{I}_{M, t, t^{-1}(u)} \circ \mathcal{J}_{\mathbf{M}, t, u}. \end{aligned} \quad (2)$$

Using (2) to substitute for $\varphi_{M'}(\lambda(v), v + \delta)$ in (1) and simplifying gives us:

$$\begin{aligned} f_v \circ \varphi_M(u, v) &= \mathcal{J}_{\mathbf{M}, t, v+\delta}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(v+\delta)}^{-1} \circ \varphi_M(\lambda(v), v + \delta) \circ \varphi_M(v, \lambda(v)) \circ \varphi_M(u, v) \\ &= \mathcal{J}_{\mathbf{M}, t, v+\delta}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(v+\delta)}^{-1} \circ \varphi_M(u, v + \delta). \end{aligned}$$

On the other hand we have, using (2) again,

$$\begin{aligned} &\varphi_{M'}(u + \delta, v + \delta) \circ f_u \\ &= \varphi_{M'}(u + \delta, v + \delta) \circ \varphi_{M'}(\lambda(u), u + \delta) \circ \mathcal{J}_{\mathbf{M}, t, \lambda(u)}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(\lambda(u))}^{-1} \circ \varphi_M(u, \lambda(u)) \\ &= \varphi_{M'}(\lambda(u), v + \delta) \circ \mathcal{J}_{\mathbf{M}, t, \lambda(u)}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(\lambda(u))}^{-1} \circ \varphi_M(u, \lambda(u)) \\ &= \mathcal{J}_{\mathbf{M}, t, v+\delta}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(v+\delta)}^{-1} \circ \varphi_M(\lambda(u), v + \delta) \circ \mathcal{I}_{M, t, t^{-1}(\lambda(u))} \\ &\quad \circ \mathcal{J}_{\mathbf{M}, t, \lambda(u)} \circ \mathcal{J}_{\mathbf{M}, t, \lambda(u)}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(\lambda(u))} \circ \varphi_M(u, \lambda(u)) \\ &= \mathcal{J}_{\mathbf{M}, t, v+\delta}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(v+\delta)}^{-1} \circ \varphi_M(\lambda(u), v + \delta) \circ \varphi_M(u, \lambda(u)) \\ &= \mathcal{J}_{\mathbf{M}, t, v+\delta}^{-1} \circ \mathcal{I}_{M, t, t^{-1}(v+\delta)}^{-1} \circ \varphi_M(u, v + \delta). \end{aligned}$$

Thus $f_v \circ \varphi_M(u, v) = \varphi_{M'}(u + \delta, v + \delta) \circ f_u$, as we wanted to show.

Finally, we need to check that $g \circ f = S(M, 2\delta)$ and $f \circ g = S(M', 2\delta)$. We perform the first verification and omit the second, since the verifications are similar.

For $u \in \mathbb{R}$,

$$\begin{aligned}
g_{u+\delta} \circ f_u &= \varphi_M(\lambda(u+\delta), u+2\delta) \circ \mathcal{I}_{M,t,t^{-1}(\lambda(u+\delta))} \circ \mathcal{J}_{\mathbf{M},t,\lambda(u+\delta)} \\
&\quad \circ \varphi_{M'}(u+\delta, \lambda(u+\delta)) \circ \varphi_{M'}(\lambda(u), u+\delta) \circ \mathcal{J}_{\mathbf{M},t,\lambda(u)}^{-1} \circ \mathcal{I}_{M,t,t^{-1}(\lambda(u))}^{-1} \circ \varphi_M(u, \lambda(u)) \\
&= \varphi_M(\lambda(u+\delta), u+2\delta) \circ \mathcal{I}_{M,t,t^{-1}(\lambda(u+\delta))} \circ \mathcal{J}_{\mathbf{M},t,\lambda(u+\delta)} \\
&\quad \circ \varphi_{M'}(\lambda(u), \lambda(u+\delta)) \circ \mathcal{J}_{\mathbf{M},t,\lambda(u)}^{-1} \circ \mathcal{I}_{M,t,t^{-1}(\lambda(u))}^{-1} \circ \varphi_M(u, \lambda(u)).
\end{aligned}$$

Using (2) once again, we have that this last expression is equal to

$$\begin{aligned}
&\varphi_M(\lambda(u+\delta), u+2\delta) \circ \mathcal{I}_{M,t,t^{-1}(\lambda(u+\delta))} \circ \mathcal{J}_{\mathbf{M},t,\lambda(u+\delta)} \circ \mathcal{J}_{\mathbf{M},t,\lambda(u+\delta)}^{-1} \circ \mathcal{I}_{M,t,t^{-1}(\lambda(u+\delta))}^{-1} \\
&\quad \circ \varphi_M(\lambda(u), \lambda(u+\delta)) \circ \mathcal{I}_{M,t,t^{-1}(\lambda(u))} \circ \mathcal{J}_{\mathbf{M},t,\lambda(u)} \circ \mathcal{J}_{\mathbf{M},t,\lambda(u)}^{-1} \circ \mathcal{I}_{M,t,t^{-1}(\lambda(u))}^{-1} \circ \varphi_M(u, \lambda(u)) \\
&= \varphi_M(\lambda(u+\delta), u+2\delta) \circ \varphi_M(\lambda(u), \lambda(u+\delta)) \circ \varphi_M(u, \lambda(u)) \\
&= \varphi_M(u, u+2\delta). \quad \square
\end{aligned}$$

Now we can complete the proof of Theorem 5.2. As mentioned above, by Theorem 5.1 it suffices to show $d_I(M, N) \leq d_B(M, N)$. Say $d_B(M, N) = \epsilon$. Choose $\delta > 0$. By Lemma 5.5, there exist well behaved modules M', N' with $d_I(M, M') \leq \delta$, $d_I(N, N') \leq \delta$. Then by Theorem 5.1, $d_B(M, M') \leq \delta$, $d_B(N, N') \leq \delta$, so by the triangle inequality, $d_B(M', N') \leq \epsilon + 2\delta$. By Lemma 5.4, $d_I(M', N') \leq \epsilon + 2\delta$. Applying the triangle inequality again, we get that $d_I(M, N) \leq \epsilon + 4\delta$. As δ may be taken to be arbitrarily small, we have $d_I(M, N) \leq \epsilon$, which completes the proof. \square

6 d_I Restricts to a Metric on Finitely Presented B_n -persistence Modules

We now show that for finitely presented B_n -modules M and N , if $d_I(M, N) = \epsilon$ then M and N are ϵ -interleaved. This implies that the restriction of d_I to finitely presented persistence modules is a metric and, as noted in Section 5, yields a converse to the algebraic stability of persistence for finitely presented B_1 -persistence modules. Theorem 6.1 will also be of some use to us in Section 11.

Theorem 6.1. *If M and N are finitely presented B_n -modules and $d_I(M, N) = \epsilon$ then M and N are ϵ -interleaved.*

Corollary 6.2. *d_I is a metric on finitely presented B_n -modules.*

Corollary 6.3 (Converse to Algebraic Stability).

- (i) *If M and N are finitely presented B_1 -persistence modules and $d_B(M, N) = \epsilon$ then M and N are ϵ -interleaved.*
- (ii) *If M and N are tame B_1 -persistence modules and $d_B(M, N) = \epsilon$ then M and N are $(\epsilon + \delta)$ -interleaved for any $\delta > 0$.*

Proof. (ii) follows directly from Theorem 5.2. (i) is immediate from that theorem and Theorem 6.1. \square

For a finitely presented B_n -persistence module M , let $U_M \subset \mathbb{R}^n$ be the set of grades of the generators and relations in a minimal presentation for M . Let $U_M^i \subset \mathbb{R}$ be the set of i^{th} coordinates of the elements of U_M .

Proof of Theorem 6.1.

Lemma 6.4. *If M is a finitely presented B_n -persistence module then for any $a \leq b \in \mathbb{R}^n$ such that $(a_i, b_i] \cap U_M^i = \emptyset$ for all i , $\varphi_M(a, b)$ is an isomorphism.*

Proof. This is straightforward; we omit the details. \square

Lemma 6.5. *If M is a finitely presented B_n -persistence module then for any $y \in \mathbb{R}^n$, there exists $r \in \mathbb{R}_{>0}$ such that $\varphi_M(y, y + r')$ is an isomorphism for all $0 \leq r' \leq r$.*

Proof. This is an immediate consequence of Lemma 6.4. \square

For a finitely presented B_n -persistence module M , let $fl_M : \mathbb{R}^n \rightarrow \prod_{i=1}^n \bar{U}_M^i$ be defined by $fl_M(a_1, \dots, a_n) = (a'_1, \dots, a'_n)$, where a'_i is the largest element of U_M^i such that $a'_i \leq a_i$, if such an element exists, and $a'_i = -\infty$ otherwise.

Lemma 6.6. *For any finitely presented B_n -module M and any $y \in \mathbb{R}^n$ with $fl_M(y) \in \mathbb{R}^n$, we have that $\varphi_M(fl_M(y), y)$ is an isomorphism.*

Proof. This too is an immediate consequence of Lemma 6.4. \square

Having stated these preliminary results, we proceed with the proof of Theorem 6.1. By Lemma 6.5 and the finiteness of U_M and U_N , there exists $\delta > 0$ such that for all $z \in U_M$, $\varphi_N(z + \epsilon, z + \epsilon + \delta)$ and $\varphi_M(z + 2\epsilon, z + 2\epsilon + 2\delta)$ are isomorphisms, and for all $z \in U_N$, $\varphi_M(z + \epsilon, z + \epsilon + \delta)$ and $\varphi_N(z + 2\epsilon, z + 2\epsilon + 2\delta)$ are isomorphisms.

By Remark 2.1, since $d_I(M, N) = \epsilon$, M and N are $(\epsilon + \delta)$ -interleaved.

Theorem 6.1 then follows from the following lemma, which will also be the key ingredient in the proof of Proposition 11.7.

Lemma 6.7. *Let M and N be finitely presented B_n -persistence modules and let $\epsilon \geq 0$ and $\delta > 0$ be such that*

1. *M and N are $\epsilon + \delta$ -interleaved,*
2. *for all $z \in U_M$, $\varphi_N(z + \epsilon, z + \epsilon + \delta)$ and $\varphi_M(z + 2\epsilon, z + 2\epsilon + 2\delta)$ are isomorphisms,*
3. *for all $z \in U_N$, $\varphi_M(z + \epsilon, z + \epsilon + \delta)$ and $\varphi_N(z + 2\epsilon, z + 2\epsilon + 2\delta)$ are isomorphisms.*

Then M and N are ϵ -interleaved.

Proof. Let $f : M \rightarrow N(\epsilon + \delta)$ and $g : N \rightarrow M(\epsilon + \delta)$ be interleaving homomorphisms.

We define ϵ -interleaving homomorphisms $\tilde{f} : M \rightarrow N(\epsilon)$ and $\tilde{g} : N \rightarrow M(\epsilon)$ via their action on homogeneous summands. First, for $z \in U_M$ define $\tilde{f}_z = \varphi_N^{-1}(z + \epsilon, z + \epsilon + \delta) \circ f_z$. Then for arbitrary $z \in \mathbb{R}^n$ such that $fl_M(z) \in \mathbb{R}^n$ define $\tilde{f}_z = \varphi_N(fl_M(z) + \epsilon, z + \epsilon) \circ \tilde{f}_{fl_M(z)} \circ$

$\varphi_M^{-1}(fl_M(z), z)$. (Note that $\varphi_M^{-1}(fl_M(z), z)$ is well defined by Lemma 6.6.) Finally, for $z \in \mathbb{R}^n$ s.t. $fl_M(z) \notin \mathbb{R}^n$, define $\tilde{f}_z = 0$. (If $fl_M(z) \notin \mathbb{R}^n$ then $M_z = 0$, so this last part of the definition is reasonable.)

Symmetrically, for $z \in U_N$ define $\tilde{g}_z = \varphi_M^{-1}(z + \epsilon, z + \epsilon + \delta) \circ g_z$. For arbitrary $z \in \mathbb{R}^n$ such that $fl_N(z) \in \mathbb{R}^n$ define $\tilde{g}_z = \varphi_M(fl_N(z) + \epsilon, z + \epsilon) \circ \tilde{g}_{fl_N(z)} \circ \varphi_N^{-1}(fl_N(z), z)$. For $z \in \mathbb{R}^n$ s.t. $fl_N(z) \notin \mathbb{R}^n$, define $\tilde{g}_z = 0$.

We need to check that \tilde{f}, \tilde{g} as thus defined are in fact morphisms. We perform the check for \tilde{f} ; the check for \tilde{g} is the same.

If $y \in \mathbb{R}^n$ is such that $fl_M(y) \notin \mathbb{R}^n$, then since $M_y = 0$, it's clear that $\tilde{f}_z \circ \varphi_M(y, z) = \varphi_N(y + \epsilon, z + \epsilon) \circ \tilde{f}_y$.

For $y \leq z \in \mathbb{R}^n$ such that $fl_M(y) \in \mathbb{R}^n$,

$$\begin{aligned}
\tilde{f}_z \circ \varphi_M(y, z) &= \varphi_N(fl_M(z) + \epsilon, z + \epsilon) \circ \tilde{f}_{fl_M(z)} \circ \varphi_M^{-1}(fl_M(z), z) \circ \varphi_M(y, z) \\
&= \varphi_N(fl_M(z) + \epsilon, z + \epsilon) \circ \varphi_N^{-1}(fl_M(z) + \epsilon, fl_M(z) + \epsilon + \delta) \circ f_{fl_M(z)} \\
&\quad \circ \varphi_M^{-1}(fl_M(z), z) \circ \varphi_M(y, z) \\
&= \varphi_N(fl_M(z) + \epsilon, z + \epsilon) \circ \varphi_N^{-1}(fl_M(z) + \epsilon, fl_M(z) + \epsilon + \delta) \circ f_{fl_M(z)} \\
&\quad \circ \varphi_M(fl_M(y), fl_M(z)) \circ \varphi_M^{-1}(fl_M(y), y) \\
&= \varphi_N(fl_M(z) + \epsilon, z + \epsilon) \circ \varphi_N^{-1}(fl_M(z) + \epsilon, fl_M(z) + \epsilon + \delta) \\
&\quad \circ \varphi_N(fl_M(y) + \epsilon + \delta, fl_M(z) + \epsilon + \delta) \circ f_{fl_M(y)} \circ \varphi_M^{-1}(fl_M(y), y) \\
&= \varphi_N(y + \epsilon, z + \epsilon) \circ \varphi_N(fl_M(y) + \epsilon, y + \epsilon) \\
&\quad \circ \varphi_N^{-1}(fl_M(y) + \epsilon, fl_M(y) + \epsilon + \delta) \circ f_{fl_M(y)} \circ \varphi_M^{-1}(fl_M(y), y) \\
&= \varphi_N(y + \epsilon, z + \epsilon) \circ \tilde{f}_y
\end{aligned}$$

as desired.

To finish the proof, we need to check that $\tilde{g} \circ \tilde{f} = S(M, 2\epsilon)$ and $\tilde{f} \circ \tilde{g} = S(N, 2\epsilon)$. We perform the first check; the second check is the same.

For $z \in \mathbb{R}^n$, if $fl_M(z) \notin \mathbb{R}^n$ then since $M_z = 0$, $\tilde{g}_{z+\epsilon} \circ \tilde{f}_z = 0 = \varphi_M(z, z + 2\epsilon)$.

To show that the result also holds for z such that $fl_M(z) \in \mathbb{R}^n$, we'll begin by verifying the result for $z \in U_M$. We'll use this special case in proving the result for arbitrary $z \in \mathbb{R}^n$ such that $fl_M(z) \in \mathbb{R}^n$.

If $z \in U_M$ then, by assumption, $\varphi_M(z + 2\epsilon, z + 2\epsilon + 2\delta)$ is an isomorphism. Thus, to show that $\tilde{g}_{z+\epsilon} \circ \tilde{f}_z = \varphi_M(z, z + 2\epsilon)$, it suffices to show that $\varphi_M(z + 2\epsilon, z + 2\epsilon + 2\delta) \circ \tilde{g}_{z+\epsilon} \circ \tilde{f}_z = \varphi_M(z, z + 2\epsilon + 2\delta)$.

For $z \in U_M$, we have

$$\begin{aligned}
\tilde{g}_{z+\epsilon} \circ \tilde{f}_z &= \varphi_M(fl_N(z + \epsilon) + \epsilon, z + 2\epsilon) \circ \tilde{g}_{fl_N(z+\epsilon)} \circ \varphi_N^{-1}(fl_N(z + \epsilon), z + \epsilon) \circ \tilde{f}_z \\
&= \varphi_M(fl_N(z + \epsilon) + \epsilon, z + 2\epsilon) \circ \varphi_M^{-1}(fl_N(z + \epsilon) + \epsilon, fl_N(z + \epsilon) + \epsilon + \delta) \circ g_{fl_N(z+\epsilon)} \\
&\quad \circ \varphi_N^{-1}(fl_N(z + \epsilon), z + \epsilon) \circ \tilde{f}_z \\
&= \varphi_M(fl_N(z + \epsilon) + \epsilon, z + 2\epsilon) \circ \varphi_M^{-1}(fl_N(z + \epsilon) + \epsilon, fl_N(z + \epsilon) + \epsilon + \delta) \circ g_{fl_N(z+\epsilon)} \\
&\quad \circ \varphi_N^{-1}(fl_N(z + \epsilon), z + \epsilon) \circ \varphi_N^{-1}(z + \epsilon, z + \epsilon + \delta) \circ f_z.
\end{aligned}$$

Thus

$$\begin{aligned}
& \varphi_M(z + 2\epsilon, z + 2\epsilon + 2\delta) \circ \tilde{g}_{z+\epsilon} \circ \tilde{f}_z \\
&= \varphi_M(z + 2\epsilon, z + 2\epsilon + \delta) \circ \varphi_M(fl_N(z + \epsilon) + \epsilon, z + 2\epsilon) \circ \varphi_M^{-1}(fl_N(z + \epsilon) + \epsilon, fl_N(z + \epsilon) + \epsilon + \delta) \\
&\quad \circ g_{fl_N(z+\epsilon)} \circ \varphi_N^{-1}(fl_N(z + \epsilon), z + \epsilon) \circ \varphi_N^{-1}(z + \epsilon, z + \epsilon + \delta) \circ f_z \\
&= \varphi_M(z + 2\epsilon + \delta, z + 2\epsilon + 2\delta) \circ \varphi_M(fl_N(z + \epsilon) + \epsilon + \delta, z + 2\epsilon + \delta) \\
&\quad \circ g_{fl_N(z+\epsilon)} \circ \varphi_N^{-1}(fl_N(z + \epsilon), z + \epsilon) \circ \varphi_N^{-1}(z + \epsilon, z + \epsilon + \delta) \circ f_z \\
&= \varphi_M(z + 2\epsilon + \delta, z + 2\epsilon + 2\delta) \circ g_{z+\epsilon} \circ \varphi_N(fl_N(z + \epsilon), z + \epsilon) \\
&\quad \circ \varphi_N^{-1}(fl_N(z + \epsilon), z + \epsilon) \circ \varphi_N^{-1}(z + \epsilon, z + \epsilon + \delta) \circ f_z \\
&= g_{z+\epsilon+\delta} \circ \varphi_N(z + \epsilon, z + \epsilon + \delta) \circ \varphi_N^{-1}(z + \epsilon, z + \epsilon + \delta) \circ f_z \\
&= g_{z+\epsilon+\delta} \circ f_z \\
&= \varphi(z, z + 2\epsilon + 2\delta)
\end{aligned}$$

as desired.

Finally, for arbitrary $z \in \mathbb{R}^n$ such that $fl_M(z) \notin \mathbb{R}^n$, we have, using that \tilde{g} is a morphism,

$$\begin{aligned}
\tilde{g}_{z+\epsilon} \circ \tilde{f}_z &= \varphi_M(fl_N(z + \epsilon) + \epsilon, z + 2\epsilon) \circ \tilde{g}_{fl_N(z+\epsilon)} \\
&\quad \circ \varphi_N^{-1}(fl_N(z + \epsilon), z + \epsilon) \circ \varphi_N(fl_M(z) + \epsilon, z + \epsilon) \circ \tilde{f}_{fl_M(z)} \circ \varphi_M^{-1}(fl_M(z), z) \\
&= \tilde{g}_{z+\epsilon} \circ \varphi_N(fl_N(z + \epsilon), z + \epsilon) \circ \varphi_N^{-1}(fl_N(z + \epsilon), z + \epsilon) \\
&\quad \circ \varphi_N(fl_M(z) + \epsilon, z + \epsilon) \circ \tilde{f}_{fl_M(z)} \circ \varphi_M^{-1}(fl_M(z), z) \\
&= \tilde{g}_{z+\epsilon} \circ \varphi_N(fl_M(z) + \epsilon, z + \epsilon) \circ \tilde{f}_{fl_M(z)} \circ \varphi_M^{-1}(fl_M(z), z) \\
&= \varphi_M(fl_M(z) + 2\epsilon, z + 2\epsilon) \circ \tilde{g}_{fl_M(z)+\epsilon} \circ \tilde{f}_{fl_M(z)} \circ \varphi_M^{-1}(fl_M(z), z) \\
&= \varphi_M(fl_M(z) + 2\epsilon, z + 2\epsilon) \circ \varphi_M(fl_M(z), fl_M(z) + 2\epsilon) \circ \varphi_M^{-1}(fl_M(z), z) \\
&= \varphi_M(z, z + 2\epsilon)
\end{aligned}$$

as we wanted. □

This completes the proof of Theorem 6.1. □

Remark 6.1. As noted in Remark 4.1, the notion of a well behaved persistence module admits a generalization to the multi-dimensional setting. An interesting question is whether Theorem 6.1 generalizes to well behaved multidimensional persistence modules; if it does, then we obtain corresponding generalizations of Corollaries 6.2 and 6.3. Our proof of Theorem 6.1 does not generalize directly.

7 Geometric Preliminaries

7.1 CW-complexes and Cellular homology

Our proof of the optimality of the interleaving distance will involve the construction of CW-complexes and the computation of their cellular homology. We now briefly review finite dimensional CW-complexes and cellular homology.

7.1.1 Definition of a Finite-dimensional CW-complex

A CW-complex is a topological space X together with some additional data of *attaching maps* specifying how X is assembled as the union of open disks of various dimensions. We quote the procedural definition of a finite-dimensional CW-complex given in [27].

Let D^i denote the unit disk in \mathbb{R}^i ; for α contained in some indexing set (which will often be implicit in our notation) let D_α^i be a copy of D^i . When i is clear from context, we will sometimes denote D_α^i simply as D_α .

A finite-dimensional CW-complex is a space X constructed in the following way:

1. Start with a discrete set X^0 , the 0-cells of X .
2. Inductively, form the i -skeleton of X^i from X^{i-1} by attaching i -cells e_α^i via maps $\sigma_\alpha : S^{i-1} \rightarrow X^{i-1}$. This means that X^i is the quotient space of $X^{i-1} \amalg_\alpha D_\alpha^i$ under the identifications $x \sim \sigma_\alpha(x)$ for $x \in \delta D_\alpha^i$. The cell e_α^i is the homeomorphic image of $D_\alpha^i - \delta D_\alpha^i$ under the quotient map.
3. $X = X^r$ for some r . We call the smallest such r the *dimension* of X .

The *characteristic map* of the cell e_α^i is the map $\Phi_\alpha : D_\alpha^i \rightarrow X$ which is the composition $D_\alpha^i \hookrightarrow X^{i-1} \amalg_\alpha D_\alpha^i \rightarrow X^i \hookrightarrow X$, where the middle map is the quotient map defining X^i .

A *subcomplex* of a CW-complex X is a closed subspace A of X which is a union of the cells of X ; those cells contained in A are taken to have the same attaching maps as they do in X .

7.1.2 Cellular Homology

We mention only what we need about cellular homology to prove Theorem 10.1. For a more complete discussion and proofs of the results stated here, see e.g. [27] or [2].

For $i \in \mathbb{Z}_{\geq 0}$, we'll let H_i denote the i^{th} singular homology functor with coefficients in the field k .

For X a CW-complex and $i \in \mathbb{N}$, let $d_i^X : H_i(X^i, X^{i-1}) \rightarrow H_{i-1}(X^{i-1}, X^{i-2})$ denote the map induced by the boundary map in the long exact sequence of the pair (X^i, X^{i-1}) . It can be checked that the d_i^X give

$$\cdots \xrightarrow{\delta_{i+1}^X} H_i(X^i, X^{i-1}) \xrightarrow{\delta_i^X} H_{i-1}(X^{i-1}, X^{i-2}) \xrightarrow{\delta_{i-1}^X} \cdots \xrightarrow{\delta_1^X} H_0(X^0) \rightarrow 0$$

the structure of a chain complex, and that the i^{th} homology vector space of this chain complex, denoted $H_i^{CW}(X)$, is isomorphic to $H_i(X)$.

It can be shown that a choice of generator for $H_i(D^i, S^{i-1}) \cong \mathbb{Z}$ induces a choice of basis for $H_i(X^i, X^{i-1})$ whose elements correspond bijectively to the i -cells of X . We now fix a choice of generator $H_i(D^i, S^{i-1})$ for each $i \in \mathbb{N}$.³ We can then think of $H_i(X^i, X^{i-1})$ as the k -vector space generated by the i -cells of X .

It follows from the equality $H_0^{CW}(X) = H_0(X)$ that in the case that X has a single 0-cell, $d_1^X = 0$.

For $i > 1$, the *cellular boundary formula* gives an explicit expression for d_i^X . To prepare for the formula, we note first that for $i \in \mathbb{N}$, the choice of generator for $H_i(D^i, S^{i-1})$ induces a

³Such a choice is induced e.g. by the standard orientation on D^i .

choice of generator a_i for $H_i(D^i/S^{i-1})$ via the quotient map $D^i \rightarrow D^i/S^{i-1}$. Also, the choice of generator for $H_{i+1}(D^{i+1}, S^i)$ induces a choice of generator b_i for $H_i(S^i)$ via the boundary map in the long exact sequence of the pair (D^{i+1}, S^i) . For each $i \in \mathbb{N}$, choose $\rho^i : D^i/S^{i-1} \rightarrow S^i$ to be any homeomorphism such that $\rho_*^i : H_i(D^i/S^{i-1}) \rightarrow H_i(S^i)$ sends a_i to b_i .

For $i \in \mathbb{N}$ and an i -cell e_β^i of X , let $(e_\beta^i)^c$ denote the compliment of e_β^i in X^i , and let $q_\beta : X^i \rightarrow X^i/(e_\beta^i)^c$ denote the quotient map. ρ^i and Φ_β induce an identification of $q_\beta(X^i)$ with S^i .

By a compactness argument [27, Section A.1], for any i -cell e_α^i the image of the attaching map σ_α of e_α^i meets only finitely many cells.

For $i > 1$, the cellular boundary formula states that

$$\delta_i^X(e_\alpha^i) = \sum_{\text{im}(\sigma_\alpha) \cup e_\beta^{i-1} \neq \emptyset} \deg(q_\beta \circ \sigma_\alpha) e_\beta^{i-1}.$$

Here, for any map $f : S^{i-1} \rightarrow S^{i-1}$, $\deg(f)$ denotes the field element $a \in k$ such that $f_* : H_{i-1}(S^{i-1}) \rightarrow H_{i-1}(S^{i-1})$ is multiplication by a .

We can endow the set of CW-complexes with the structure of a category by taking $\text{hom}(X, Y)$ for CW-complexes X, Y to be the set of continuous maps $f : X \rightarrow Y$ such that $f(X^i) \subset Y^i$ for all i . We call maps $f \in \text{hom}(X, Y)$ *cellular maps*. It can be shown that a cellular map f induces a map $H_i^{CW}(f) : H_i^{CW}(X) \rightarrow H_i^{CW}(Y)$ in such a way that H_i^{CW} becomes a functor.

Further, there exists a natural isomorphism [28] $\kappa : H_i^{CW} \rightarrow \bar{H}_i$, where \bar{H}_i is the restriction of H_i to the category of CW-complexes.

7.2 Filtrations

Fix $n \in \mathbb{N}$.

Define an **n -filtration** X to be a collection of topological spaces $\{X_a\}_{a \in \mathbb{R}^n}$, together with a collection of continuous injections $\{\phi_X(a, b) : X_a \rightarrow X_b\}_{a \leq b \in \mathbb{R}^n}$ such that if $a \leq b \leq c \in \mathbb{R}^n$ then $\phi_X(b, c) \circ \phi_X(a, b) = \phi_X(a, c)$. Given two n -filtrations X and Y , we define a morphism f from X to Y to be a collection of continuous functions $\{f_a\}_{a \in \mathbb{R}^n} : X_a \rightarrow Y_a$ such that for all $a \leq b \in \mathbb{R}^n$, $f_b \circ \phi_X(a, b) = \phi_Y(a, b) \circ f_a$. This definition of morphism gives the n -filtrations the structure of a category. Let $n\text{-filt}$ denote this category.

Define a **cellular n -filtration** to be a collection of CW-complexes $\{X_a\}_{a \in \mathbb{R}^n}$, together with inclusions of subcomplexes $\{\phi_X(a, b) : X_a \rightarrow X_b\}_{a \leq b \in \mathbb{R}^n}$. Given two cellular n -filtrations X and Y , we define a morphism f from X to Y to be a collection of cellular maps $\{f_a\}_{a \in \mathbb{R}^n} : X_a \rightarrow Y_a$ such that for all $a \leq b \in \mathbb{R}^n$, $f_b \circ \phi_X(a, b) = \phi_Y(a, b) \circ f_a$. This definition of morphism gives the cellular n -filtrations the structure of a category.

Simplicial n -filtrations can be defined analogously.

7.3 Multidimensional Persistent Homology

The multidimensional persistent homology functor H_i is a generalization of the ordinary homology functor with field coefficients to the setting where the source is an n -filtration and the

target is a B_n -module. We first present a definition of the singular multidimensional persistent homology functor; we introduce cellular multidimensional persistent homology below.

7.3.1 Singular Multidimensional Persistent Homology

For a topological space X and $j \in \mathbb{Z}_{\geq 0}$, let $C_j(X)$ denote the j^{th} singular chain module of X , with coefficients in k . For X, Y topological spaces and $f : X \rightarrow Y$ a continuous map, let $f^\# : C_j(X) \rightarrow C_j(Y)$ denote the map induced by f .

For X an n -filtration, define $C_j(X)$, the j^{th} singular chain module of X , as the B_n -persistence module for which $C_j(X)_u = C_j(X_u)$ for all $u \in \mathbb{R}^n$ and for which $\varphi_{C_j(X)}(u, v) = \phi_X^\#(u, v)$. Note that for any $j \in \mathbb{Z}_{\geq 0}$, the collection of boundary maps $\{\delta_j : C_j(X_u) \rightarrow C_{j-1}(X_u)\}_{u \in \mathbb{R}^n}$ induces a boundary map $\delta_j : C_j(X) \rightarrow C_{j-1}(X)$. These boundary maps give $\{C_j(X)\}_{j \geq 0}$ the structure of a chain complex. We define the $H_j(X)$, the j^{th} persistent homology module of X , to be the j^{th} homology module of this complex. For X and Y two n -filtrations, a morphism $f \in \text{hom}(X, Y)$ induces in the obvious way a morphism $H_j(f) : H_j(X) \rightarrow H_j(Y)$, making $H_j : n\text{-filt} \rightarrow B_n\text{-mod}$ a functor.

7.3.2 Cellular Multidimensional Persistent Homology

A construction analogous to the one above, with cellular chain complexes used in place of singular chain complexes, yields a definition of the cellular multidimensional persistent homology of cellular n -filtrations. For a cellular filtration X , let $C_j^{CW}(X)$ denote the j^{th} cellular chain module of X , let $\delta_j^{CW} : C_j^{CW}(X) \rightarrow C_{j-1}^{CW}(X)$ denote the j^{th} cellular chain map of X , and let $H_j^{CW}(X)$ denote the j^{th} cellular multidimensional persistence module of X .

Remark 7.1. It follows from the naturality of the isomorphisms between singular and cellular homology that the singular and cellular multidimensional persistent homology modules of a cellular n -filtration are isomorphic.

7.4 Functors from Geometric Categories to Categories of n -filtrations

In a typical application of persistent homology, one has some geometric⁴ category of interest and a functor F from that category to $n\text{-filt}$; one then studies and works with the functors $H_i \circ F$. This composite functor is also generally referred to an i^{th} persistent homology functor, and there are a number of different such i^{th} persistent homology functors with different sources, each determined by a different choice of the functor F .

We next examine several examples of functors $F : C \rightarrow n\text{-filt}$, where C is some geometric category. These will prove important to us in Section 9, where we will formulate definitions of optimality of pseudometrics on B_n -persistence modules in a way which depends on a choice of F and a metric on $\text{obj}^*(C)$.

In the following examples, we omit the specification of the action of these functors on morphisms. This should be clear from context.

⁴We are using the word “geometric” here in an informal—and rather broad—sense.

Example 7.1. Sublevelset Filtrations

Let C^S be the category defined as follows:

1. Objects of C^S are pairs (X, f) , where X is a topological space and $f = (f_1, \dots, f_n) : X \rightarrow \mathbb{R}^n$ is a function.
2. If $(X, f), (X', f') \in \text{obj}(C^S(n))$, then we define $\text{hom}_{C^S}((X, f), (X', f'))$ to be the set of functions $\gamma : X \rightarrow X'$ such that $f(x) \geq f'(\gamma(x))$ for all $x \in X$.

For X a topological space, let C_X^S denote the subcategory of C^S whose objects are pairs of the form (X, f) .

For $a = (a_1, \dots, a_n) \in \mathbb{R}^n$, let $X_a = \cap_{i=1}^n f_i^{-1}((-\infty, a_i])$. If $a, b \in \mathbb{R}^n$ with $a \leq b$, then $X_a \subset X_b$. Thus the collection of subsets $\{X_r\}_{r \in \mathbb{R}^n}$ is an n -filtration. We call the functor which maps the pair (X, f) to this n -filtration the *sublevelset filtration functor*, and denote it F^S .

This functor has previously been considered in [10].

Example 7.2. Sublevelset-Offset Filtrations

Let C^{SO} be the category defined as follows:

1. Objects of C^{SO} are triplets (X, d, f) , where (X, d) is a metric space and $f : X \rightarrow \mathbb{R}^n$ is a function.
2. If $(X, d, f), (X', d', f') \in \text{obj}(C^{SO})$, then we define $\text{hom}_{C^{SO}}((X, d, f), (X', d', f'))$ to be the set of functions $\gamma : X \rightarrow X'$ such that $f(x) \geq f'(\gamma(x))$ for all $x \in X$ and $d(x, y) \geq d'(\gamma(x), \gamma(y))$ for all $x, y \in X$.

For X a topological space, let C_X^{SO} denote the subcategory of C^{SO} consisting of triplets of the form (X, d, f) .

For $a \in \mathbb{R}, b \in \mathbb{R}^n$ let X_a be defined as in Example 7.1, and let $X_{(a,b)} = \{x \in X | d(x, X_a) \leq b\}$. If $(a, b) \leq (c, d) \in \mathbb{R}^n \times \mathbb{R} = \mathbb{R}^{n+1}$, then $X_{(a,b)} \subset X_{(c,d)}$. Thus the collection of subsets $\{X_{(a,b)}\}_{(a,b) \in \mathbb{R}^n \times \mathbb{R}}$ is an $(n+1)$ -filtration. We call the functor which maps the triple (X, d, f) to this $(n+1)$ -filtration the *sublevelset-offset filtration functor*, and denote it F^{SO} .

It appears that the functor F^{SO} has not previously been considered in the computational topology literature. However, it is an easy common generalization of the 1-dimensional filtrations considered for example in [17] and [15].

By taking the metric information into account, the functor F^{SO} encodes more nuanced information about the topography of the functions on metric spaces than the functor F^S does. For instance, when (X, d) is \mathbb{R}^2 with the Euclidean metric, and $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a function, $H_i \circ F^{SO}(X, d, f)$ encodes information about the width and shape of topographical features of the graph of f that $H_i \circ F^S(X, f)$ would not capture.

Example 7.3. Sublevelset-Rips Filtrations

Let C^{SR} be the subcategory of C^{SO} whose objects are triples (X, d, f) , where X is a finite metric space.

Given a finite metric space (X, d) , and $b \in \mathbb{R}_{\geq 0}$, let $R(X, d, b)$, denote the Rips complex of (X, d) with parameter b [12]. If $b < 0$, let $R(X, d, b) = R(X, d, 0)$.

Given $(X, d, f) \in \text{obj}(C^{SR})$, let X_a be defined as in Example 7.1 and let d_a be the restriction of d to $X_a \times X_a$. Now for $a, b \in \mathbb{R}$, let $X_{(a,b)} = R(X_a, d_a, b)$.

If $(a, b) \leq (c, d) \in \mathbb{R}^n \times \mathbb{R}$ then $X_{(a,b)} \subset X_{(c,d)}$, so the collection of subsets $\{X_{(a,b)}\}_{(a,b) \in \mathbb{R}^n \times \mathbb{R}}$ is an $(n+1)$ -filtration. We call the functor which maps the triple (X, d, f) to this $(n+1)$ -filtration the *sublevelset-Rips filtration functor*, and denote it F^{SR} .

This functor has previously been considered (in the case $n = 1$) in [8].

7.5 Metrics on Geometric Categories

We now define metrics on the isomorphism classes of objects of the geometric categories we defined in the last section.

Let X be a topological space.

7.5.1 A Metric on $\text{obj}^*(C_X^S)$

We define a metric d_X^S on $\text{obj}^*(C_X^S)$ by taking $d_X^S((X, f_1), (X, f_2)) = \sup_{x \in X} \|f_1(x) - f_2(x)\|_\infty$.

7.5.2 A Metric on $\text{obj}^*(C_X^{SO})$

We define a metric d_X^{SO} on $\text{obj}^*(C_X^{SO})$ by taking

$$d_X^S((X, d_1, f_1), (X, d_2, f_2)) = \max(\sup_{x \in X} \|f_1(x) - f_2(x)\|_\infty, \sup_{x_1, x_2 \in X} |d_1(x_1, x_2) - d_2(x_1, x_2)|).$$

7.5.3 A Metric on $\text{obj}^*(C^{SR})$

Generalizing in a mild way a definition of [12], we define a metric d^{SR} on $\text{obj}^*(C^{SR})$. (In fact, this definition extends to the subcategory of C^{SO} whose objects are the triplets (X, d_X, f_X) with X compact, but we won't need the extra generality here.) For $f_X, f_Y \equiv 0$, $d^{SR}((X, d_X, f_X), (X, d_X, f_X))$ will be equal to the Gromov-Hausdorff metric [12].

To define d^{SR} , we need some preliminary definitions and notation. Define a *correspondence* between two sets X and Y to be a subset $C \in X \times Y$ such that $\forall x \in X, \exists y \in Y$ such $(x, y) \in C$, and $\forall y \in Y, \exists x \in X$ s.t. $(x, y) \in C$. Let $C(X, Y)$ denote the set of correspondences between X and Y .

For $(X, d_X, f_X), (Y, d_Y, f_Y) \in C^{SR}$, define $\Gamma_{X,Y} : X \times Y \times X \times Y \rightarrow \mathbb{R}_{\geq 0}$ by

$$\Gamma_{X,Y}(x, y, x', y') = |d_X(x, x') - d_Y(y, y')|.$$

For $C \in C(X, Y)$, define Γ_C as $\sup_{(x,y),(x',y') \in C} \Gamma_{X,Y}(x, y, x', y')$, and define $|f_X - f_Y|_C$ to be $\sup_{(x,y) \in C} \|f_X(x) - f_Y(y)\|_\infty$. Informally, Γ_C is the maximum distortion of the metrics under the correspondence C , and $|f_X - f_Y|_C$ is the maximum distortion of the functions under C .

Now define we define d^{SR} by taking

$$d^{SR}((X, d_X, f_X), (Y, d_Y, f_Y)) = \inf_{C \in C(X,Y)} \max(\frac{1}{2}\Gamma_C, |f_X - f_Y|_C).$$

7.6 Stability Results for Ordinary Persistence

There are two main geometric stability results for ordinary persistence in the literature. (Though see also the generalization [5]) Each is a consequence of the algebraic stability of persistence [11].

Theorem 7.1 (1-D Stability Result for C_X^S [11]). *For any $i \in \mathbb{Z}_{\geq 0}$, topological space X , and functions $f_1, f_2 : X \rightarrow \mathbb{R}$ such that $H_i \circ F^S(X, f_1)$ and $H_i \circ F^S(X, f_2)$ are tame,*

$$d_B(H_i \circ F^S(X, f_1), H_i \circ F^S(X, f_2)) \leq d_X^S((X, f_1), (X, f_2)).$$

For a 2-D filtration F , let $\text{diag}(F)$ denote the 1-D filtration for which $\text{diag}(F)_a = F_{(a,a)}$.

Theorem 7.2 (1-D Stability Result for C^{SR} [12]). *For finite metric spaces $(X, d_X), (Y, d_Y)$ and functions $f_X : X \rightarrow \mathbb{R}, f_Y : Y \rightarrow \mathbb{R}$,*

$$d_B(H_i \circ \text{diag}(F^{SR}(X, d_X, f_X)), H_i \circ \text{diag}(F^{SR}(Y, d_Y, f_Y))) \leq d^{SR}((X, d_X, f_X), (Y, d_Y, f_Y)).$$

We'll see in Section 8 that both of these results admit generalizations to the setting of multidimensional persistence in terms of the interleaving metric.

8 Stability Properties of the Interleaving Distance

In this section, we observe that multidimensional persistent homology is stable with respect to the interleaving distance in three senses analogous to those in which ordinary persistent homology is known to be stable. As noted in the introduction, there is not much mathematical work to do here. Nevertheless, these observations are significant not only because they show that the interleaving distance is in certain respects a well behaved distance, but also because stability is closely related to the optimality of distances on persistence modules as we define it in Section 9; insofar as we wish to understand the optimality properties of the interleaving distance, the stability properties of the interleaving distance are of interest.

As before, fix $n \in \mathbb{N}$.

8.1 Multidimensional Persistence Stability Result 1

Theorem 8.1. *For any topological space X and pairs $(X, f_1), (X, f_2) \in \text{obj}(C_X^S)$ we have, for any $i \in \mathbb{Z}_{\geq 0}$,*

$$d_I(H_i \circ F^S(X, f_1), H_i \circ F^S(X, f_2)) \leq d_X^S((X, f_1), (X, f_2)).$$

The case $n = 1$ is Theorem 7.1.

Proof. Let $d_X^S((X, f_1), (X, f_2)) = \epsilon$. Then for any $u \in \mathbb{R}^n$, $F^S(X, f_1)_u \subset F^S(X, f_2)_{u+\epsilon}$ and $F^S(X, f_2)_u \subset F^S(X, f_1)_{u+\epsilon}$. The images of these inclusions under the i^{th} singular homology functor define ϵ -interleaving morphisms between $H_i \circ F^S(X, f_1)$ and $H_i \circ F^S(X, f_2)$. Thus $d_I(H_i \circ F^S(X, f_1), H_i \circ F^S(X, f_2)) \leq \epsilon$ as needed. \square

8.2 Multidimensional Persistence Stability Result 2

Theorem 8.2. *For any topological space X and triples $(X, d_1, f_1), (X, d_2, f_2) \in \text{obj}(C_X^{SO})$ we have, for any $i \in \mathbb{Z}_{\geq 0}$,*

$$d_I(H_i \circ F^{SO}(X, d_1, f_1), H_i \circ F^{SO}(X, d_2, f_2)) \leq d_X^{SO}((X, d_1, f_1), (X, d_2, f_2)).$$

Proof. This is similar to proof of the previous result. $d_X^{SO}((X, d_1, f_1), (X, d_2, f_2)) = \epsilon$. Then for any $u \in \mathbb{R}^n$ and $r \in \mathbb{R}$, $F^{SO}(X, f_1)_{(u, r)} \subset F^{SO}(X, f_2)_{(u+\epsilon, r+\epsilon)}$ and $F^{SO}(X, f_2)_{(u, r)} \subset F^{SO}(X, f_1)_{(u+\epsilon, r+\epsilon)}$. The result now follows via the same argument given in the proof of Theorem 8.1. \square

8.3 Multidimensional Persistence Stability Result 3

Theorem 8.3. *For $(X, d_X, f_X), (Y, d_Y, f_Y) \in \text{obj}(C^{SR})$ we have, for any $i \in \mathbb{Z}_{\geq 0}$,*

$$d_I(H_i \circ F^{SR}(X, d_X, f_X), H_i \circ F^{SR}(Y, d_Y, f_Y)) \leq d^{SR}((X, d_X, f_X), (Y, d_Y, f_Y)).$$

The proof of this is a very minor modification of the argument given in [12] to prove Theorem 7.2.

Note that Theorem 8.3 implies Theorem 7.2: When $n = 1$, if $H_i \circ F^{SR}(X, d_X, f_X)$ and $H_i \circ F^{SR}(Y, d_Y, f_Y)$ are ϵ -interleaved, for $\epsilon \in \mathbb{R}_{\geq 0}$, then $H_i \circ \text{diag}(F^{SR}(X, d_X, f_X))$ and $H_i \circ \text{diag}(F^{SR}(Y, d_Y, f_Y))$ are ϵ -interleaved.

9 Optimal Pseudometrics

In this section we introduce a relative notion of optimality of pseudometrics on persistence modules and their discrete invariants. This relative notion of optimality is quite general and specializes to a number of different notions of optimality of pseudometrics of interest in the context of multidimensional persistence.

We also present some first theoretical observations about optimal pseudometrics. We'll exploit these in Section 10 to prove our optimality result for the interleaving distance.

9.1 A General Definition of Optimal Pseudometrics

Let Y be a set. We define a **relative structure** on Y to be a triple $\mathcal{R} = (\mathcal{T}, X_{\mathcal{T}}, f_{\mathcal{T}})$, where \mathcal{T} is a set, $X_{\mathcal{T}} = \{(X_s, d_s)\}_{s \in \mathcal{T}}$ is a collection of pseudometric spaces indexed by \mathcal{T} , and $f_{\mathcal{T}} = \{f_s : X_s \rightarrow Y\}_{s \in \mathcal{T}}$ is a collection of functions. Let $\text{im}(f_{\mathcal{T}}) = \cup_{s \in \mathcal{T}} \text{im}(f_s)$.

If Y is a set and $\mathcal{R} = (\mathcal{T}, X_{\mathcal{T}}, f_{\mathcal{T}})$ is a relative structure on Y , we say a semi-pseudometric d on Y is **\mathcal{R} -stable** if for every $s \in \mathcal{T}$ and $x_1, x_2 \in X_s$ we have $d(f_s(x_1), f_s(x_2)) \leq d_s(x_1, x_2)$.

We say a pseudometric d on Y is **\mathcal{R} -optimal** if d is \mathcal{R} -stable, and for every other \mathcal{R} -stable pseudometric d' on Y , we have $d'(y_1, y_2) \leq d(y_1, y_2)$ for all $y_1, y_2 \in \text{im}(f_{\mathcal{T}})$.

The following lemma is immediate, but important to understanding our definition of \mathcal{R} -optimality.

Lemma 9.1. *An \mathcal{R} -stable pseudometric d is \mathcal{R} -optimal iff for any other \mathcal{R} -stable pseudometric d' , $s \in \mathcal{T}$, and $x_1, x_2 \in X_s$,*

$$|d_s(x_1, x_2) - d(f_s(x_1), f_s(x_2))| \leq |d_s(x_1, x_2) - d'(f_s(x_1), f_s(x_2))|.$$

Note that if an \mathcal{R} -optimal pseudometric d on Y exists, its restriction to $\text{im}(f_{\mathcal{T}}) \times \text{im}(f_{\mathcal{T}})$ is unique.

9.2 Examples

Here we give several examples of sets Y and relative structures $(\mathcal{T}, X_{\mathcal{T}}, f_{\mathcal{T}})$ on Y for which it would be interesting or useful from the standpoint of the theory and application of multidimensional persistent homology to identify an \mathcal{R} -optimal pseudometric. In Section 10 we will focus exclusively on the relative structures of Examples 9.1 and 9.2; we leave it to future work to investigate in detail the optimality of pseudometrics with respect to the relative structures of Examples 9.3-9.6.

For examples 9.1-9.5, let $Y = \text{obj}^*(B_n\text{-mod})$.

Example 9.1. Let $\mathcal{R}_1 = (\mathcal{T}, X_{\mathcal{T}}, f_{\mathcal{T}})$, where \mathcal{T} is the set of pairs $\{(T, i) | T \text{ is a topological space, } i \in \mathbb{Z}_{\geq 0}\}$, $X_{(T, i)} = \text{obj}^*(C_T^S)$ for $(T, z) \in \mathcal{T}$, $d_{(T, i)} = d_T^S$, and $f_{(T, i)}$ is given by $f_{(T, i)}(T, g) = H_i \circ F^S(T, g)$.

Remark 9.1. By Lemma 9.1, an \mathcal{R}_1 -stable pseudometric d is \mathcal{R}_1 -optimal iff for any other \mathcal{R}_1 -stable pseudometric d' , any $(T, i) \in \mathcal{T}$, and any $(T, g_1), (T, g_2) \in X_{(T, i)}$,

$$\begin{aligned} & |d_T^S((T, g_1), (T, g_2)) - d(H_i \circ F^S(T, g_1), H_i \circ F^S(T, g_2))| \\ & \leq |d_T^S((T, g_1), (T, g_2)) - d'(H_i \circ F^S(T, g_1), H_i \circ F^S(T, g_2))|. \end{aligned}$$

This says that an \mathcal{R}_1 -optimal pseudometric is one for which the L_{∞} distance between any two functions defined on the same topological space is preserved under the multidimensional persistent homology functor as faithfully as is possible for any choice of \mathcal{R}_1 -stable pseudometric on $\text{obj}^*(B_n\text{-mod})$.

For the relative structures \mathcal{R} in the rest of the examples below, \mathcal{R} -optimality has an interpretation by way of Lemma 9.1 analogous to that of Remark 9.1.

Example 9.2. For $i \in \mathbb{Z}_{\geq 0}$, let $\mathcal{R}_{1, i} = (\mathcal{T}, X_{\mathcal{T}}, f_{\mathcal{T}})$, where \mathcal{T} is the set of topological spaces, $X_T = \text{obj}^*(C_T^S)$ for $T \in \mathcal{T}$, $d_T = d_T^S$, and f_T is given by $f_T(T, g) = H_i \circ F^S(T, g)$.

Since the definitions of \mathcal{R}_1 and $\mathcal{R}_{1, i}$ are similar, one might expect that there's a relationship between \mathcal{R}_1 -optimality and $\mathcal{R}_{1, i}$ -optimality. Corollary 10.2 establishes such a relationship in the cases $k = \mathbb{Q}$ and $k = \mathbb{Z}/p\mathbb{Z}$ for some prime p .

Example 9.3. Let \mathcal{T} be the set of pairs $\{(T, i) | T \text{ is a topological space, } i \in \mathbb{Z}_{\geq 0}\}$, $X_{(T, i)} = \text{obj}^*(C_T^{SO})$ for $(T, z) \in \mathcal{T}$, $d_{(T, i)} = d_T^{SO}$, and $f_{(T, i)}$ be given by $f_{(T, i)}(T, d, g) = H_i \circ F^{SO}(T, d, g)$.

Example 9.4. Let \mathcal{T} be the singleton set $\{s\}$. Let $X_s = \text{obj}^*(C^{SR})$, $d_s = d^{SR}$, and f_s be given by $f_s(T, d, g) = H_i \circ F^{SR}(T, d, g)$.

Example 9.5. We can present variants of examples 9.3 and 9.4 where we only consider homology in a single dimension, in the same way we did for Example 9.1 in Example 9.2.

Example 9.6. Let $W : \text{obj}^*(B_n\text{-mod}) \rightarrow Y$ be a discrete invariant [8] with values in a set Y , let $(\mathcal{T}, X_{\mathcal{T}}, f'_{\mathcal{T}})$ be any relative structure on $\text{obj}^*(B_n\text{-mod})$, and let $f_{\mathcal{T}} = W \circ f'_{\mathcal{T}}$. Then $(\mathcal{T}, X_{\mathcal{T}}, f_{\mathcal{T}})$ is a relative structure on W .

For example, we can take W to be the rank invariant [8] and $(\mathcal{T}, X_{\mathcal{T}}, f'_{\mathcal{T}})$ to be the relative structure of Example 9.1.

9.3 Induced Semi-Pseudometrics and a Condition for the Existence of Optimal Pseudometrics

We'll see here that a relative structure $\mathcal{R} = (\mathcal{T}, X_{\mathcal{T}}, f_{\mathcal{T}})$ on a set Y induces a semi-pseudometric $d_{\mathcal{R}}$ on $\text{im}(f_{\mathcal{T}})$ with a nice property.

For $y_1, y_2 \in \text{im}(f_{\mathcal{T}})$, let

$$A(y_1, y_2) = \{(s, x_1, x_2) | s \in \mathcal{T}, x_1 \in X_s, x_2 \in X_s, f_s(x_1) = y_1, f_s(x_2) = y_2\}.$$

Now define $d_{\mathcal{R}}(y_1, y_2) = \inf_{(s, x_1, x_2) \in A(y_1, y_2)} d_s(x_1, x_2)$. $d_{\mathcal{R}}$ is an \mathcal{R} -stable semi-pseudometric. In general it needn't satisfy the triangle inequality. However, if \mathcal{T} is a singleton set, as for instance in Example 9.4, then $d_{\mathcal{R}}$ does satisfy the triangle inequality and is a pseudometric.

Lemma 9.2. *for any \mathcal{R} -stable pseudometric d on $\text{im}(f_{\mathcal{T}})$, $d \leq d_{\mathcal{R}}$. In particular, if $d_{\mathcal{R}}$ is a pseudometric, it is \mathcal{R} -optimal.*

Proof. Let d be an \mathcal{R} -stable pseudometric on $\text{im}(f_{\mathcal{T}})$. Since d is \mathcal{R} -stable, $d(y_1, y_2) \leq d_s(x_1, x_2)$ for all $(s, x_1, x_2) \in A(y_1, y_2)$. Thus $d(y_1, y_2) \leq \inf_{(s, x_1, x_2) \in A(y_1, y_2)} d_s(x_1, x_2) = d_{\mathcal{R}}(y_1, y_2)$. \square

It's easy to see that a pseudometric d on $\text{im}(f_{\mathcal{T}})$ can be extended (non-uniquely) to a metric on Y ; if d is \mathcal{R} -optimal then, by definition, any extension to a pseudometric on Y is as well. Thus by the Lemma 9.2, if $d_{\mathcal{R}}$ is a pseudometric, then an \mathcal{R} -optimal pseudometric exists on Y ; its restriction to $\text{im}(f_{\mathcal{T}})$ is $d_{\mathcal{R}}$. It follows, for example, that for \mathcal{R} as in Example 9.4 an \mathcal{R} -optimal metric exists.

In Section 10, we'll show that for $Y = \text{obj}^*(B_n\text{-mod})$, $i \in \mathbb{N}$, and $k = \mathbb{Q}$ or $k = \mathbb{Z}/p\mathbb{Z}$ for some prime p , the restriction of d_I to the domain of $d_{\mathcal{R}_{1,i}}$ is equal to $d_{\mathcal{R}_{1,i}}$, so that d_I is $\mathcal{R}_{1,i}$ -optimal. It will follow easily that d_I is also \mathcal{R}_1 -optimal.

10 Optimality of The Interleaving Distance (Relative to Sublevelset Persistence)

10.1 Summary of Results

The central result of this section is the following theorem:

Theorem 10.1. *For $k = \mathbb{Q}$ or $\mathbb{Z}/p\mathbb{Z}$ for some prime p , and $i \in \mathbb{N}$, d_I is $\mathcal{R}_{1,i}$ -optimal.*

This theorem also yields the following weaker optimality result, which has aesthetic advantage of not depending in its formulation on a choice of homology dimension.

Corollary 10.2. *For $k = \mathbb{Q}$ or $\mathbb{Z}/p\mathbb{Z}$ for some prime p , d_I is \mathcal{R}_1 -optimal.*

We also present variants of Theorems 10.1 and 10.2 for well behaved B_1 -persistence modules which hold for arbitrary fields k ; these are given as Theorem 10.6. Our analogue of Theorem 10.1 for well behaved B_1 -persistence modules holds for $i \in \mathbb{Z}_{\geq 0}$.

As an intermediate step in our proof of Theorem 10.1, we prove Theorem 10.7, which we believe to be of independent interest. Theorem 10.7 gives a condition equivalent to the existence of ϵ -interleaving homomorphisms between two persistence modules. It expresses transparently the sense in which ϵ -interleaved persistence modules are algebraically similar.

10.2 Proofs

Proof of Theorem 10.1. Fix $k = \mathbb{Q}$ or $k = \mathbb{Z}/p\mathbb{Z}$ for some prime p . For $i \in \mathbb{Z}_{\geq 0}$, Lemma 9.2 implies that to prove that d_I is $\mathcal{R}_{1,i}$ -optimal, it's enough to show that the restriction of d_I to the domain of $d_{\mathcal{R}_{1,i}}$ is equal to $d_{\mathcal{R}_{1,i}}$. We show that for $i \in \mathbb{N}$, this follows from the following proposition.

Proposition 10.3. *Let $k = \mathbb{Q}$ or $k = \mathbb{Z}/p\mathbb{Z}$ for some prime p . If $i \in \mathbb{N}$, and M and N are ϵ -interleaved B_n -modules, then there exists a CW-complex X and continuous functions $\gamma_M, \gamma_N : X \rightarrow \mathbb{R}^n$ such that $H_i \circ F^S(X, \gamma_M) \cong M$, $H_i \circ F^S(X, \gamma_N) \cong N$, and $\|\gamma_M - \gamma_N\|_\infty = \epsilon$.*

Sections 10.4-10.8 are devoted to the proof of Proposition 10.3.

Corollary 10.4. *For $k = \mathbb{Q}$ or $\mathbb{Z}/p\mathbb{Z}$ for some prime p , for every B_n -module M and $i \in \mathbb{N}$, there exists a CW-complex X and a continuous function $\gamma : X \rightarrow \mathbb{R}^n$ such that $H_i \circ F^S(X, \gamma) \cong M$.*

To see that Proposition 10.3 implies that d_I is equal to $d_{\mathcal{R}_{1,i}}$ on their common domain, let M and N be two B_n -persistence modules such that $d_I(M, N) = \epsilon$. Choose $\delta > 0$. Then M and N are $(\epsilon + \delta)$ -interleaved. By the proposition, there exists a topological space X and maps $\gamma_M, \gamma_N : X \rightarrow \mathbb{R}^n$ such that $H_i \circ F^S(X, \gamma_M) \cong M$, $H_i \circ F^S(X, \gamma_N) \cong N$, and $\|\gamma_M - \gamma_N\|_\infty = \epsilon + \delta$. Thus by stability, $d_{\mathcal{R}_{1,i}}(M, N) \leq \epsilon + \delta$. Since this holds for all $\delta > 0$, $d_{\mathcal{R}_{1,i}}(M, N) \leq \epsilon$. By Lemma 9.2, $d_{\mathcal{R}_{1,i}}(M, N) = \epsilon$. \square

Note that the extension of Theorem 10.1 to the case $i = 0$ would follow by the same argument from the following conjectural extension of Proposition 10.3 to the case $i = 0$.

Conjecture 10.1. *Let $k = \mathbb{Q}$ or $k = \mathbb{Z}/p\mathbb{Z}$ for some prime p . Let M and N be ϵ -interleaved B_n -modules such that for some topological spaces X_M, X_N and functions $f_M : X_M \rightarrow \mathbb{R}^n$, $f_N : X_N \rightarrow \mathbb{R}^n$, $H_0 \circ F^S(X_M, f_M) \cong M$ and $H_0 \circ F^S(X_N, f_N) \cong N$. Then there exists a CW-complex X and continuous functions $\gamma_M, \gamma_N : X \rightarrow \mathbb{R}^n$ such that $H_0 \circ F^S(X, \gamma_M) \cong M$, $H_0 \circ F^S(X, \gamma_N) \cong N$, and $\|\gamma_M - \gamma_N\|_\infty = \epsilon$.*

We do not prove this conjecture here.

Proof of Corollary 10.2. Fix any $i \in \mathbb{N}$. Write $\mathcal{R}_1 = (S, X_S, f_S)$, $\mathcal{R}_{1,i} = (S', X_{S'}, f_{S'})$. By Theorem 8.1, d_I is \mathcal{R}_1 -stable. Further, any \mathcal{R}_1 -stable pseudometric d' on $\text{obj}^*(B_n\text{-mod})$ is $\mathcal{R}_{1,i}$ -stable. By Corollary 10.4, when $k = \mathbb{Q}$ or $\mathbb{Z}/p\mathbb{Z}$ for some prime p , $\text{im}(f_{S'}) = \text{im}(f_S) = \text{obj}^*(B_n\text{-mod})$. Thus, if d' is any \mathcal{R}_1 stable metric, $d'(M, N) \leq d_I(M, N)$ for any $M, N \in \text{im}(f_S)$ by the $\mathcal{R}_{1,i}$ -optimality of d_I . Thus d_I is \mathcal{R}_1 -optimal. \square

10.3 Optimality of the Bottleneck Distance

Before proceeding with the proof of Proposition 10.3 we present optimality results for the bottleneck distance.

Let Y denote the set of isomorphism classes of tame B_1 -persistence modules. Since the bottleneck distance d_B is only defined between elements of Y , formulating statements about the optimality of d_B requires that we consider a relative structure on Y rather than on $\text{obj}^*(B_1\text{-mod})$. We can (in the obvious way) define restrictions \mathcal{R}_2 and $\mathcal{R}_{2,i}$ of the relative structures \mathcal{R}_1 and $\mathcal{R}_{1,i}$ to relative structures on Y .

Then, given Proposition 10.3, the proofs of Theorem 10.1 and Corollary 10.2 adapt to give that for $k = \mathbb{Q}$ or $k = \mathbb{Z}/p\mathbb{Z}$ for p a prime, d_B is \mathcal{R}_2 -optimal, and for any $i \in \mathbb{N}$, d_B is also $\mathcal{R}_{2,i}$ -optimal.

Moreover, if we further restrict our attention to well behaved B_1 -persistence modules then we can prove an optimality result for arbitrary fields k and persistent homology of any dimension.

The key is the following variant of Proposition 10.3:

Proposition 10.5. *Let k be any field and let $i \in \mathbb{Z}_{\geq 0}$. Let M and N be well behaved B_1 -persistence modules with $d_B(M, N) = \epsilon$. If $i = 0$, assume further that for some topological spaces X_M, X_N and functions $f_M : X_M \rightarrow \mathbb{R}$, $f_N : X_N \rightarrow \mathbb{R}$, $H_0 \circ F^S(X_M, f_M) \cong M$ and $H_0 \circ F^S(X_N, f_N) \cong N$. Then for any $\delta > 0$, there exists a CW-complex X and continuous functions $\gamma_M, \gamma_N : X \rightarrow \mathbb{R}$ such that $H_i \circ F^S(X, \gamma_M) \cong M$, $H_i \circ F^S(X, \gamma_N) \cong N$, and $\|\gamma_M - \gamma_N\|_\infty \leq \epsilon + \delta$.*

Proof. An easy constructive proof, similar on a high level to our proof of Proposition 10.3 below but much simpler, follows from the definition of d_B and the structure theorem for well behaved persistence modules (Theorem 4.5). We leave the details to the reader. \square

Now let Y' denote the set of isomorphism classes of well behaved B_1 -persistence modules. Let \mathcal{R}_3 and $\mathcal{R}_{3,i}$ denote the restrictions of the relative structures \mathcal{R}_1 and $\mathcal{R}_{1,i}$ to relative structures on Y' .

Our optimality result for the bottleneck distance on well behaved persistence modules is the following:

Theorem 10.6. *For any field k and $i \in \mathbb{Z}_{\geq 0}$, d_B is $\mathcal{R}_{3,i}$ -optimal. Further, d_B is \mathcal{R}_3 -optimal.*

Proof. Using Proposition 10.5 in place of Proposition 10.3, the proofs of Theorem 10.1 and Corollary 10.2 carry over with only minor modifications to give the result. \square

Remark 10.1. In the case of 0-D ordinary persistence, our Theorem 10.6 implies a slight weakening of the optimality result of [19]. It is easy to check that in the geometric context considered in [19], the persistent homology modules obtained must satisfy property 2 in the definition of a well behaved B_1 -persistence module, but they needn't satisfy property 1. To strengthen our optimality result to a full generalization of the result of [19], we'd want a generalization of Theorem 4.5 to all B_1 -persistence modules satisfying property 2. We presume that this can be obtained via a slight strengthening of Theorem 3.2, but we do not pursue the details of this here.

10.4 Proof of Proposition 10.3, Part 1: A Characterization of ϵ -interleaved Pairs of Modules

The rest of Section 10 is devoted to the proof of Proposition 10.3.

The main step in the proof of Proposition 10.3 is the proof of Theorem 10.7 below, which as we have noted, gives a condition equivalent to the existence of ϵ -interleaving homomorphisms between two persistence modules. The reader should think of the less trivial direction of Theorem 10.7 as an algebraic analogue of Proposition 10.3.

Remark 10.2. Before stating the theorem, we point out that for any n -graded set G and $\epsilon \geq 0$, the homomorphism $S(\langle G(-\epsilon) \rangle, \epsilon) : \langle G(-\epsilon) \rangle \rightarrow \langle G \rangle$ is injective, and so gives an identification $\langle G(-\epsilon) \rangle$ with a subset of $\langle G \rangle$.

More generally, if G_1 and G_2 are n -graded sets, we obtain in the obvious way an identification of $\langle G_1, G_2(-\epsilon) \rangle$ with a subset of $\langle G_1, G_2 \rangle$.

Theorem 10.7. *Let M and N be B_n -persistence modules. For any $\epsilon \in \mathbb{R}_{\geq 0}$, M and N are ϵ -interleaved if and only if there exist n -graded sets $\mathcal{W}_1, \mathcal{W}_2$ and sets $\mathcal{Y}_1, \mathcal{Y}_2 \subset \langle \mathcal{W}_1, \mathcal{W}_2 \rangle$ such that $\mathcal{Y}_1 \in \langle \mathcal{W}_1, \mathcal{W}_2(-\epsilon) \rangle$, $\mathcal{Y}_2 \in \langle \mathcal{W}_1(-\epsilon), \mathcal{W}_2 \rangle$,*

$$\begin{aligned} M &\cong \langle \mathcal{W}_1, \mathcal{W}_2(-\epsilon) | \mathcal{Y}_1, \mathcal{Y}_2(-\epsilon) \rangle \\ N &\cong \langle \mathcal{W}_1(-\epsilon), \mathcal{W}_2 | \mathcal{Y}_1(-\epsilon), \mathcal{Y}_2 \rangle. \end{aligned}$$

If M and N are finitely presented, then $\mathcal{W}_1, \mathcal{W}_2, \mathcal{Y}_1, \mathcal{Y}_2$ can be taken to be finite.

Proof of Theorem 10.7. It's easy to see that if there exist n -graded sets $\mathcal{W}_1, \mathcal{W}_2$ and sets $\mathcal{Y}_1, \mathcal{Y}_2 \subset \langle \mathcal{W}_1, \mathcal{W}_2 \rangle$ as in the statement of the theorem then M and N are ϵ -interleaved.

To prove the converse, we lift to free covers of M and N a construction presented in the proof of [11, Lemma 4.6]. [11, Lemma 4.6] was stated only for B_1 -persistence modules, but the result and its proof generalize immediately to B_n -persistence modules.

Let $f : M \rightarrow N(-\epsilon)$, $g : N \rightarrow M(-\epsilon)$ be ϵ -interleaving homomorphisms.

Upon generalizing to B_n -persistence modules, the proof of Lemma [11, Lemma 4.6] yields the following result as a special case:

Lemma 10.8. *Let $\gamma_1 : M(-2\epsilon) \rightarrow M \oplus N(-\epsilon)$ be given by $\gamma_1(y) = (S(M(-2\epsilon), 2\epsilon)(y), -f(y))$. Let $\gamma_2 : N(-\epsilon) \rightarrow M \oplus N(-\epsilon)$ be given by $\gamma_2(y) = (-g(y), y)$. Let $R \subset M \oplus N(-\epsilon)$ be the submodule generated by $\text{im}(\gamma_1) \cup \text{im}(\gamma_2)$. Then $M \cong (M \oplus N(-\epsilon))/R$.*

For convenience's sake, we reprove Lemma 10.8 here.

Proof. Let $\iota : M \rightarrow M \oplus N(-\epsilon)$ denote the inclusion, and let $\zeta : M \oplus N(-\epsilon) \rightarrow M \oplus N(-\epsilon)/R$ denote the quotient. We'll show that $\zeta \circ \iota$ is an isomorphism. For any $(y_M, y_N) \in M \oplus N(-\epsilon)$, $(-g(y_N), y_N) \in R$, so $\zeta \circ \iota(g(y_N)) = (0, y_N) + R$. Therefore $\zeta \circ \iota(g(y_N) + y_M) = (y_M, y_N) + R$. Hence $\zeta \circ \iota$ is surjective.

$\zeta \circ \iota$ is injective iff $\iota(M) \cap R = 0$. It's clear that $M \cap \text{im}(\gamma_2) = 0$. Thus to show that $\zeta \circ \iota$ is injective it's enough to show that $\text{im}(\gamma_1) \subset \text{im}(\gamma_2)$. If $y \in M(-2\epsilon)$, then since $S(M(-2\epsilon), 2\epsilon)(y) = g \circ f(y)$, $(S(M(-2\epsilon), 2\epsilon)(y), -f(y)) = (g \circ f(y), -f(y)) = \gamma_2(-f(y))$. Thus $\text{im}(\gamma_1) \subset \text{im}(\gamma_2)$ and so $\zeta \circ \iota$ is injective.

Thus $\zeta \circ \iota$ is an isomorphism. \square

Now let $\langle G_M | R_M \rangle$ be a presentation for M and let $\langle G_N | R_N \rangle$ be a presentation for N . Without loss of generality we may assume $M = \langle G_M \rangle / \langle R_M \rangle$ and $N = \langle G_N \rangle / \langle R_N \rangle$. Let $\rho_M : \langle G_M \rangle \rightarrow M$, $\rho_N : \langle G_N \rangle \rightarrow N$ denote the quotient maps. Then $(\langle G_M \rangle, \rho_M)$ and $(\langle G_N \rangle, \rho_N)$ are free covers for M and N .

Let $\tilde{f} : \langle G_M \rangle \rightarrow \langle G_N(\epsilon) \rangle$ be a lift of f and let $\tilde{g} : \langle G_N \rangle \rightarrow \langle G_M(\epsilon) \rangle$ be a lift of g .

Let $R_{M,N} = \{y - \tilde{f}(y)\}_{y \in G_M(-\epsilon)}$ and let $R_{N,M} = \{y - \tilde{g}(y)\}_{y \in G_N(-\epsilon)}$. Note that $R_{M,N}$ is a homogeneous subset of $\langle G_M(-\epsilon), G_N \rangle$ and $R_{N,M}$ is a homogeneous subset of $\langle G_M, G_N(-\epsilon) \rangle$.

Let

$$\begin{aligned} P_M &= \langle G_M, G_N(-\epsilon) | R_M, R_N(-\epsilon), R_{M,N}(-\epsilon), R_{N,M} \rangle, \\ P_N &= \langle G_M(-\epsilon), G_N | R_M(-\epsilon), R_N, R_{M,N}, R_{N,M}(-\epsilon) \rangle. \end{aligned}$$

$R_{M,N}(-\epsilon)$ lies in $\langle G_M(-2\epsilon), G_N(-\epsilon) \rangle$. By Remark 10.2, the map $S(G_M(-2\epsilon), 2\epsilon)$ identifies $\langle G_M(-2\epsilon) \rangle$ with a subset of $\langle G_M \rangle$. Thus P_M is well defined. By an analogous observation, P_N is also well defined.

We claim that P_M is a presentation for M and P_N is a presentation for N . We'll prove that P_M is a presentation for M ; The proof that P_N is a presentation for N is identical.

Let

$$\begin{aligned} F &= \langle G_M, G_N(-\epsilon) \rangle, \\ K &= \langle R_M, R_N(-\epsilon), R_{M,N}(-\epsilon), R_{N,M} \rangle \\ K' &= \langle R_M, R_N(-\epsilon) \rangle. \end{aligned}$$

Let $p : F \rightarrow F/K$ denote the quotient map. Clearly, we may identify F/K' with $M \oplus N(-\epsilon)$. We'll check that under this identification, p maps $\langle R_{M,N}(-\epsilon) \rangle$ surjectively to $\text{im}(\gamma_1)$ and $\langle R_{N,M} \rangle$ surjectively to $\text{im}(\gamma_2)$. Thus $K/K' = R$. It follows that P_M is a presentation for M by Lemma 10.8 and the third isomorphism theorem for modules [22].

We first check that $\langle p(R_{M,N}(-\epsilon)) \rangle = \text{im}(\gamma_1)$. Viewing $R_{M,N}(-\epsilon)$ as a subset of $\langle G_M, G_N(-\epsilon) \rangle$, $R_{M,N}(-\epsilon) = \{S(G_M(-2\epsilon), 2\epsilon)(y) - \tilde{f}(y)\}_{y \in G_M(-2\epsilon)}$. $S(G_M(-2\epsilon), 2\epsilon)$ is a lift of $S(M(-2\epsilon), 2\epsilon)$ and \tilde{f} is a lift of f , so for any $y \in G_M(-2\epsilon)$,

$$p(S(G_M(-2\epsilon), 2\epsilon)(y) - \tilde{f}(y)) = (S(M(-2\epsilon), 2\epsilon)(\rho_M(y)), -f(\rho_M(y))) = \gamma_1(\rho_M(y)).$$

Thus $p(R_{M,N}(-\epsilon)) \subset \text{im}(\gamma_1)$. Since G_M generates $\langle G_M \rangle$ and ρ_M is surjective, we have that $p(\langle R_{M,N}(-\epsilon) \rangle) = \text{im}(\gamma_1)$.

The check that $\langle p(R_{N,M}) \rangle = \text{im}(\gamma_2)$ is similar to the above verification that $\langle p(R_{M,N}(-\epsilon)) \rangle = \text{im}(\gamma_1)$, but simpler. $R_{N,M} = \{y - \tilde{g}(y)\}_{y \in G_N(-\epsilon)}$. \tilde{g} is a lift of g so for any $y \in G_N(-\epsilon)$,

$$p(y - \tilde{g}(y)) = (-g(\rho_N(y)), \rho_N(y)) = \gamma_2(\rho_N(y)).$$

Thus $p(R_{N,M}) \subset \text{im}(\gamma_2)$. Since G_N generates $\langle G_N \rangle$ and ρ_N is surjective, we have that $p(\langle R_{N,M} \rangle) = \text{im}(\gamma_2)$.

This completes the verification that P_M is a presentation for M .

Now, taking $\mathcal{W}_1 = G_M$, $\mathcal{W}_2 = G_N$, $\mathcal{Y}_1 = R_M \cup R_{N,M}$, and $\mathcal{Y}_2 = R_N \cup R_{M,N}$ gives the first statement of Theorem 10.7. If M and N are finitely presented then $G_M, G_N, R_M, R_N, R_{M,N}$, and $R_{N,M}$ can all be taken to be finite; the second statement of Theorem 10.7 follows. \square

10.5 Proof of Proposition 10.3, Part 2: Constructing the CW-complex

Given $\mathcal{W}_1, \mathcal{W}_2, \mathcal{Y}_1, \mathcal{Y}_2$ as in the statement of Theorem 10.7, we now proceed with the proof of Proposition 10.3, beginning with the construction of the CW-complex X whose existence is posited by the proposition.

Write $\mathcal{W} = \mathcal{W}_1 \cup \mathcal{W}_2$ and $\mathcal{Y} = \mathcal{Y}_1 \cup \mathcal{Y}_2$. If $M = N = 0$, Proposition 10.3 clearly holds, so we may assume without loss of generality that M and N are not both trivial. Under this assumption, $\mathcal{W} \neq \emptyset$.

For an n -graded set S , let $GCD(S) = (v_1, \dots, v_n)$, where $v_i = \inf_{s \in S} gr(s)_i$. In general, some of the components of $GCD(\mathcal{W})$ may be equal to $-\infty$. To simplify our exposition, we first present the remainder of the proof of Proposition 10.3 in the special case that $GCD(\mathcal{W}) \in \mathbb{R}^n$. The adaptation of our proof to the general case is reasonably straightforward; we outline this adaptation in Section 10.8.

We'll define X so that

1. X has a single 0-cell B .
2. X has an i -cell e_w^i for each $w \in \mathcal{W}$.
3. X has an $(i+1)$ -cell e_y^{i+1} for each $y \in \mathcal{Y}$.

For such X , the attaching map for each i -cell e_w^i must be the constant map to B . To define X , then, we need only to specify the attaching map $\sigma_y : S^i \rightarrow X^i$ for each $y \in \mathcal{Y}$.

We do this for $k = \mathbb{Q}$, and leave to the reader the easy adaptation of the construction (and its use in the remainder of the proof of Proposition 10.3) to the case $k = \mathbb{Z}/p\mathbb{Z}$.

For any $y \in \mathcal{Y}$, we may choose a finite set $\mathcal{W}_y \subset \mathcal{W}$ such that $gr(w) \leq gr(y)$ for each $w \in \mathcal{W}_y$, and

$$y = \sum_{w \in \mathcal{W}_y} a'_{wy} \varphi_{\langle \mathcal{W} \rangle}(gr(w), gr(y))(w) \quad (3)$$

for some $a'_{wy} \in \mathbb{Q}$. There's an integer z such that for each a'_{wy} in the sum, $za'_{wy} \in \mathbb{Z}$. Let $a_{wy} = za'_{wy}$. For $w \notin \mathcal{W}_y$, define $a_{wy} = 0$.

Lemma 10.9. *There exists a choice of attaching map $\sigma_y : S^i \rightarrow X^i$ for each $y \in \mathcal{Y}$ such that the CW-complex X constructed via these attaching maps has $\delta_{i+1}^X(e_y^{i+1}) = \sum_{w \in \mathcal{W}} a_{wy} e_w^i$ for all $y \in \mathcal{Y}$.*

Proof. Let ρ^i be as defined in Section 7.1.2. For each $w \in \mathcal{W}$, ρ^i and the characteristic map Φ_w induce an identification of $\text{im}(\Phi_w)$ with an i -sphere S_w^i . We have that $(X^i, B) = \bigwedge_{w \in \mathcal{W}} (S_w^i, B)$. Choose a basepoint $o \in S^i$ and for each $w \in \mathcal{W}_y$, let $\sigma_{wy} : (S^i, o) \rightarrow (S_w^i, B)$ be a based map of degree a_{wy} . $[\sigma_{wy}] \in \pi_i(X^i, B)$, where $\pi_i(X^i, B)$ denotes the i^{th} homotopy group of X^i with basepoint B .

Order the elements of \mathcal{W}_y arbitrarily and call them w_1, \dots, w_l . Let $\sigma_y : (S^i, o) \rightarrow (X^i, B)$ be a map in $[\sigma_{w_1y}] \cdot [\sigma_{w_2y}] \cdot \dots \cdot [\sigma_{w_ly}] \in \pi_i(X^i, B)$. Then for any $w \in \mathcal{W}$, $q_w \circ \sigma_y$ is a map of degree a_{wy} . (See Section 7.1.2 for the definition of q_w). By the definition of δ_{i+1}^X given in Section 7.1.2, the lemma now follows. \square

We choose the attaching maps σ_y so that δ_{i+1}^X has the property specified in Lemma 10.9.

10.6 Proof of Proposition 10.3, Part 3: Defining γ_M and γ_N

Having defined the CW-complex X , we next define $\gamma_M, \gamma_N : X \rightarrow \mathbb{R}^n$.

Let $\tilde{X} = \{B\} \amalg_{w \in \mathcal{W}} D_w^i \amalg_{y \in \mathcal{Y}} D_y^{i+1}$.

X is the quotient of \tilde{X} under the equivalence relation generated by the attaching maps of the cells of X . Let $\pi : \tilde{X} \rightarrow X$ denote the quotient map. For a topological space A , let $\mathfrak{C}(A, \mathbb{R}^n)$, denote the space of continuous functions from A to \mathbb{R}^n . The map $\tilde{\cdot} : \mathfrak{C}(X, \mathbb{R}^n) \rightarrow \mathfrak{C}(\tilde{X}, \mathbb{R}^n)$ defined by $\tilde{f}(x) = f(\pi(x))$ is a bijective correspondence between elements of $\mathfrak{C}(X, \mathbb{R}^n)$ and elements of $\mathfrak{C}(\tilde{X}, \mathbb{R}^n)$ which are constant on equivalence classes.

In what follows, we'll define γ_M and γ_N by specifying their lifts $\tilde{\gamma}_M, \tilde{\gamma}_N$.

We'll take each of our functions $\tilde{\gamma}_M, \tilde{\gamma}_N$ to have the property that for each disk of \tilde{X} , the restriction of the function to any *radial line segment* (i.e. a line segment from the origin of the disk to the boundary of the disk) is linear. Given this assumption, to specify each function it is enough to specify its values on the origins of each disk of \tilde{X} . We now do this.

For any $i \in \mathbb{N}$ and any unit disk D in \mathbb{R}^i , let $O(D)$ denote the origin of D .

- $\tilde{\gamma}_M(B) = \tilde{\gamma}_N(B) = GCD(\mathcal{W})$;
- For $x \in \mathcal{W}_1 \cup \mathcal{Y}_1$, $\tilde{\gamma}_M(O(D_x)) = gr(x)$;
- For $x \in \mathcal{W}_2 \cup \mathcal{Y}_2$, $\tilde{\gamma}_M(O(D_x)) = gr(x(-\epsilon))$.
- For $x \in \mathcal{W}_1 \cup \mathcal{Y}_1$, $\tilde{\gamma}_N(O(D_x)) = gr(x(-\epsilon))$;
- For $x \in \mathcal{W}_2 \cup \mathcal{Y}_2$, $\tilde{\gamma}_N(O(D_x)) = gr(x)$.

Lemma 10.10. $\|\gamma_M - \gamma_N\|_\infty = \epsilon$.

Proof. Assume that for a disk D of \tilde{X} , $|\tilde{\gamma}_M(a) - \tilde{\gamma}_N(a)| \leq \epsilon$ for all $a \in \delta D$, and that $|\tilde{\gamma}_M(O(D)) - \tilde{\gamma}_N(O(D))| = \epsilon$. We'll show that then $|\tilde{\gamma}_M(a) - \tilde{\gamma}_N(a)| \leq \epsilon$ for all $x \in D$. Applying this result once gives that the lemma holds on the restriction of γ_M, γ_N to X^i . Applying the result a second time gives that the lemma holds on all of X .

To show that $|\tilde{\gamma}_M(a) - \tilde{\gamma}_N(a)| \leq \epsilon$, let x be a point in D and write $a = tO(D) + (1-t)b$ for some $b \in \delta D$, and $0 \leq t \leq 1$. Since the restrictions of $\tilde{\gamma}_M$ and $\tilde{\gamma}_N$ to any radial line segment from $O(D)$ to δD are linear, we have that $\tilde{\gamma}_M(a) = t\tilde{\gamma}_M(O(D)) + (1-t)\tilde{\gamma}_M(b)$, and $\tilde{\gamma}_N(a) = t\tilde{\gamma}_N(O(D)) + (1-t)\tilde{\gamma}_N(b)$. Thus $|\tilde{\gamma}_M(a) - \tilde{\gamma}_N(a)| \leq t|\tilde{\gamma}_M(O(D)) - \tilde{\gamma}_N(O(D))| + (1-t)|\tilde{\gamma}_M(b) - \tilde{\gamma}_N(b)| \leq t\epsilon + (1-t)\epsilon = \epsilon$ as needed. \square

10.7 Proof of Proposition 10.3, Part 4: Showing that $H_i \circ F^S(X, \gamma_M) \cong M$, $H_i \circ F^S(X, \gamma_N) \cong N$

Now it remains to show that $H_i \circ F^S(X, \gamma_M) \cong M$, $H_i \circ F^S(X, \gamma_N) \cong N$. We'll show that $H_i \circ F^S(X, \gamma_M) \cong M$; the argument that $H_i \circ F^S(X, \gamma_N) \cong N$ is essentially same.

For $a \in \mathbb{R}^n$, let \mathcal{F}_a denote the subcomplex of X consisting of only those cells e such that $\gamma_M(O(D(e))) \leq a$, where in this expression $D(e)$ is the disk of \tilde{X} whose interior maps to e under π . $\{\mathcal{F}_a\}_{a \in \mathbb{R}^n}$ defines a cellular filtration, which we'll denote \mathcal{F} . Let $X_a = F^S(X, \gamma_N)_a$. It's easy to see that \mathcal{F}_a is a deformation retract of X_a . Further, the inclusions of each $\mathcal{F}_a \hookrightarrow X_a$ define a morphism χ of filtrations; this morphism of filtrations maps under H_i to a morphism $H_i(\chi) : H_i(\mathcal{F}) \rightarrow H_i(F^S(X, \gamma_M))$ of B_n -persistence modules whose maps $H_i(\chi)_a : H_i(\mathcal{F}_a) \rightarrow H_i(X_a)$ are isomorphisms. Any homomorphism of B_n -persistence modules whose action on each homogeneous summand is a vector space isomorphism must be an isomorphism of B_n -persistence modules, so $H_i(\chi)$ is an isomorphism. Thus, to prove that $H_i \circ F^S(X, \gamma_M) \cong M$, it's enough to show that $H_i(\mathcal{F}) \cong M$.

By Remark 7.1, $H_i(\mathcal{F}) \cong H_i^{CW}(\mathcal{F})$.

Note that \mathcal{F} has the property that each cell e of X has a unique minimal grade of appearance $gr_{\mathcal{F}}(e)$ in \mathcal{F} . Since each cell has a unique minimal grade of appearance, for any $j \in \mathbb{Z}_{\geq 0}$, $C_j^{CW}(\mathcal{F})$ is free:

$$C_j^{CW}(\mathcal{F}) = \oplus_{e^j \subset X_a \text{ } j\text{-cell}} B_n(-gr_{\mathcal{F}}(e^j)).$$

The usual identification of j -cells of X with a basis for $C_j^{CW}(X)$ extends in the obvious way to an identification of the j -cells of X with a basis for $C_j^{CW}(\mathcal{F})$.

Moreover, the boundary homomorphism $\delta_{i+1}^X : C_{i+1}^{CW}(X) \rightarrow C_i^{CW}(X)$ and the boundary homomorphism $\delta_{i+1}^{\mathcal{F}} : C_{i+1}^{CW}(\mathcal{F}) \rightarrow C_i^{CW}(\mathcal{F})$ are related in a simple way:

Lemma 10.11. $\delta_{i+1}^{\mathcal{F}}(e_y^{i+1}) = \sum_{w \in \mathcal{W}} a_{wy} \varphi_{C_i^{CW}(\mathcal{F})}(gr_{\mathcal{F}}(e_w^i), gr_{\mathcal{F}}(e_y^{i+1}))(e_w^i)$.

Proof. Recall that we constructed X in such a way that for any $y \in \mathcal{Y}$, $\delta_{i+1}^X(e_y^{i+1}) = \sum_{w \in \mathcal{W}} a_{wy} e_w^i$. The result follows in a routine way from this expression for $\delta_{i+1}^X(e_y^{i+1})$ and the definition of the boundary map $\delta_{i+1}^{\mathcal{F}}$. \square

Now note that we have $\delta_i^{\mathcal{F}} = 0$. If $i \neq 1$ this is follows from the fact that \mathcal{F} has no $i-1$ cells. If $i = 1$, it is still true because of the isomorphism between cellular and singular persistent homology: we must have $C_0^{CW}(\mathcal{F}) \cong B_n(-gr_{\mathcal{F}}(B)) \cong H_0(\mathcal{F}) \cong H_0^{CW}(\mathcal{F})$, so $\delta_1 = 0$.

Therefore $H_i^{CW}(\mathcal{F}) = C_i^{CW}(\mathcal{F})/\text{im}(\delta_{i+1}^{\mathcal{F}})$.

The bijection which sends $w \in \mathcal{W}$ to the cell e_w^i induces an isomorphism $\Lambda : \langle \mathcal{W}_1 \cup \mathcal{W}_2(-\epsilon) \rangle \rightarrow C_i^{CW}(\mathcal{F})$.

By the expression (3) for y in terms of a'_{wy} given in Section 10.5, for $y \in \mathcal{Y}_1 \cup \mathcal{Y}_2(-\epsilon)$,

$$\Lambda(y) = \sum_{w \in \mathcal{W}} a'_{wy} \varphi_{C_i^{CW}(\mathcal{F})}(gr_{\mathcal{F}}(e_w^i), gr_{\mathcal{F}}(e_y^{i+1}))(e_w^i).$$

Thus

$$\begin{aligned} \Lambda(\langle \mathcal{Y}_1 \cup \mathcal{Y}_2(-\epsilon) \rangle) &= \langle \{ \sum_{w \in \mathcal{W}} a'_{wy} \varphi_{C_i^{CW}(\mathcal{F})}(gr_{\mathcal{F}}(e_w^i), gr_{\mathcal{F}}(e_y^{i+1}))(e_w^i) \}_{y \in \mathcal{Y}} \rangle \\ &= \langle \{ \sum_{w \in \mathcal{W}} a_{wy} \varphi_{C_i^{CW}(\mathcal{F})}(gr_{\mathcal{F}}(e_w^i), gr_{\mathcal{F}}(e_y^{i+1}))(e_w^i) \}_{y \in \mathcal{Y}} \rangle \\ &= \text{im}(\delta_{i+1}^{\mathcal{F}}) \end{aligned}$$

by Lemma 10.11. Λ therefore descends to an isomorphism between $C_i^{CW}(\mathcal{F})/\text{im}(\delta_{i+1}^{\mathcal{F}})$ and $\langle \mathcal{W}_1, \mathcal{W}_2(-\epsilon) \rangle / \langle \mathcal{Y}_1, \mathcal{Y}_2(-\epsilon) \rangle$. This shows that $H_i^{CW}(\mathcal{F}) = M$ and thus completes the proof of Proposition 10.3 in the special case that $GCD(\mathcal{W}) \in \mathbb{R}^n$.

10.8 Proof of Proposition 10.3, Part 5: Adapting the Proof to the Case That $GCD(\mathcal{W}) \notin \mathbb{R}^n$

As noted above, the proof of Proposition 10.3 we have given for the special case that $GCD(\mathcal{W}) \in \mathbb{R}^n$ adapts readily to a proof for the general case. We now outline this adaptation, leaving to the reader the straightforward details.

Let X' be the standard CW-complex structure on \mathbb{R} . That is, we take each $z \in \mathbb{Z}$ to be a 0-cell in X' , and for each $z \in \mathbb{Z}$, we take the interval $(z, z+1)$ to be a 1-cell in X' .

To carry out the proof for the general case, we first present a modified version of the construction of the CW-complex X given in Section 10.5. In this modified version, we construct the CW-complex X so that

1. X' is a subcomplex of X .
2. X has an i -cell e_w^i for each $w \in \mathcal{W}$.
3. X has an $(i+1)$ -cell e_y^{i+1} for each $y \in \mathcal{Y}$.
4. As a set, $X = X' \amalg \{e_w^i\}_{w \in \mathcal{W}} \amalg \{e_y^{i+1}\}_{y \in \mathcal{Y}}$.

For $r = (r_1, \dots, r_n) \in \mathbb{R}^n$, let $\lfloor r \rfloor = \max\{z \in \mathbb{Z} \mid z \leq r_j \text{ for } 1 \leq j \leq n\}$. For all $w \in \mathcal{W}$, let the attaching map of e_w^i be the constant map to the 0-cell $\lfloor gr(w) \rfloor \in X'$.

This defines X^i . To complete the construction of X , it remains only to specify the attaching map $\sigma_y : S^i \rightarrow X^i$ for each $y \in \mathcal{Y}$.

Let a_{wy} be defined as in Section 10.5. An analogue of Lemma 10.9 holds in our present setting and admits a similar proof. Invoking this result, for each $y \in \mathcal{Y}$ we choose $\sigma_y : S^i \rightarrow X^i$ such that the CW-complex X constructed via these attaching maps satisfies $\delta_{i+1}^X(e_y^{i+1}) = \sum_{w \in \mathcal{W}} a_{wy} e_w^i$ for each $y \in \mathcal{Y}$.

For $r \in X' = \mathbb{R}$, we define $\gamma_M(r) = \gamma_N(r) = r\vec{1}$. As in the special case, we define the values of $\tilde{\gamma}_M$ and $\tilde{\gamma}_N$ at the origin of the disks D_w^i and D_y^{i+1} as follows:

- For $x \in \mathcal{W}_1 \cup \mathcal{Y}_1$, $\tilde{\gamma}_M(O(D_x)) = gr(x)$;

- For $x \in \mathcal{W}_2 \cup \mathcal{Y}_2$, $\tilde{\gamma}_M(O(D_x)) = gr(x(-\epsilon))$.
- For $x \in \mathcal{W}_1 \cup \mathcal{Y}_1$, $\tilde{\gamma}_N(O(D_x)) = gr(x(-\epsilon))$;
- For $x \in \mathcal{W}_2 \cup \mathcal{Y}_2$, $\tilde{\gamma}_N(O(D_x)) = gr(x)$.

As in the special case, we require the restriction of $\tilde{\gamma}_M$ and $\tilde{\gamma}_N$ to radial line segments of the disks D_w^i and D_y^{i+1} to be linear. This completes the specification of γ_M and γ_N .

The argument of Lemma 10.10 shows that $\|\gamma_M - \gamma_N\|_\infty = \epsilon$, and the argument of Section 10.7 adapts in a straightforward way to show that $H_i \circ F^S(X, \gamma_M) \cong M$ and $H_i \circ F^S(X, \gamma_N) \cong N$. \square

11 Reducing the Evaluation of d_I to Deciding Solvability of Quadratics

Let $\mathcal{MQ}(k)$ denote the set of multivariate systems of quadratic equations over the field k .

Fix $n \in \mathbb{N}$ and Let M and N be finitely presented B_n -persistence modules. Let q be the total number of generators and relations in a minimal presentation for M and in a minimal presentation for N . We show in this section that given minimal presentations for M and N , for any $\epsilon > 0$ deciding whether M and N are ϵ -interleaved is equivalent to deciding the solvability of an instance of $\mathcal{MQ}(k)$ with $O(q^2)$ unknowns and $O(q^2)$ equations.

We also show that d_I must be equal to one of the elements of an order $O(q^2)$ subset of $\mathbb{R}_{\geq 0}$ defined in terms of the grades of generators and relations of M and N . Thus, by Theorem 6.1, by searching through these values we can compute d_I by deciding whether M and N are ϵ -interleaved for $O(\log q)$ values of ϵ . That is, we can compute $d_I(M, N)$ by deciding the solvability of $O(\log q)$ instances of $\mathcal{MQ}(k)$.

If e.g. k is a field of prime order, a standard algorithm based on Gröbner bases determines the solvability of systems in $\mathcal{MQ}(k)$. $\mathcal{MQ}(k)$ is NP-complete, however, and this algorithm is for general instances of $\mathcal{MQ}(k)$ prohibitively inefficient. We leave it to future work to investigate the complexity and tractability in practice of deciding the solvability of systems in $\mathcal{MQ}(k)$ arising from our reduction.

In practice, we are interested in computing the interleaving distance between the simplicial persistent homology modules of two simplicial n -filtrations. To apply the reduction presented here to this problem, we need a way of computing a presentation of the multidimensional persistent homology module of a simplicial n -filtration; strictly speaking, our reduction does not require that the presentations of our modules be minimal. However, to minimize the number and size of the quadratic systems we need to consider in computing the interleaving distance via this reduction, we do want the presentations we compute to be minimal.

We will address the problem of computing a minimal presentation of the simplicial persistent homology module of a simplicial n -filtration in a companion paper.

11.1 Linear Algebraic Representations of Homogeneous Elements and Morphisms of Free B_n -persistence Modules

11.1.1 Representing Homogeneous Elements of Free B_n -persistence Modules as Vectors

Given a finitely generated free B_n -persistence module F and an (ordered) basis $B = b_1, \dots, b_l$ for F , we can represent a homogeneous element $v \in F$ as a pair $([v, B], \text{gr}(v))$ where $[v, B] \in k^l$ is a vector: if $v = \sum_{i: \text{gr}(v) \geq \text{gr}(b_i)} a_i \varphi_F(\text{gr}(b_i), \text{gr}(v))(b_i)$, with each $a_i \in k$, then for $1 \leq i \leq l$ we define

$$[v, B]_i = \begin{cases} a_i & \text{if } \text{gr}(v) \geq \text{gr}(b_i) \\ 0 & \text{otherwise.} \end{cases}$$

Remark 11.1. Note that for $1 \leq i \leq l$, $[b_i, B] = \mathbf{e}_i$, where \mathbf{e}_i denotes the i^{th} standard basis vector in k^l .

If $V \subset F$ is a set, we define $[V, B] = \{[v, B] | v \in V\}$.

11.1.2 Representing Morphisms of B_n -persistence Modules as Matrices

Given finitely generated B_n -persistence modules F and F' and (ordered) bases $B = b_1, \dots, b_l$ and $B' = b'_1, \dots, b'_m$ for F and F' respectively, let $\text{Mat}_k(B, B')$ denote the set of $m \times l$ matrices A with entries in k such that $A_{ij} = 0$ whenever $\text{gr}(b_j) < \text{gr}(b'_i)$.

We can represent a morphism $f \in \text{hom}(F, F')$ as a matrix $[f, B, B'] \in \text{Mat}_k(B, B')$, where if $f(b_j) = \sum_{i: \text{gr}(b_j) \geq \text{gr}(b'_i)} a_{ij} \varphi_{F'}(\text{gr}(b'_i), \text{gr}(b_j))(b'_i)$, with each $a_{ij} \in k$, then

$$[f, B, B']_{ij} = \begin{cases} a_{ij} & \text{if } \text{gr}(b_j) \geq \text{gr}(b'_i) \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 11.1. *The map $[\cdot, B, B'] : \text{hom}(F, F') \rightarrow \text{Mat}_k(B, B')$ is a bijection.*

Proof. The proof is straightforward. □

Note also the following additional properties of these matrix representations of morphisms between free B_n -modules:

Lemma 11.2. *Let F, F', F'' be free B_n -persistence modules with ordered bases B, B', B'' .*

- (i) *If $f_1, f_2 \in \text{hom}(F, F')$ then $[f_1 + f_2, B, B'] = [f_1, B, B'] + [f_2, B, B']$,*
- (ii) *If $f_1 \in \text{hom}(F, F')$, $f_2 \in \text{hom}(F', F'')$ then $[f_2 \circ f_1, B, B''] = [f_2, B', B''] [f_1, B, B']$,*
- (iii) *For any $\epsilon \geq 0$, $[S(F, \epsilon), B, B(\epsilon)] = I_{|B|}$, where for $m \in \mathbb{N}$, I_m denotes the $m \times m$ identity matrix.*

Proof. The proof of each of these results is straightforward. □

For a graded set W and $u \in \mathbb{R}^n$, let $W^u = \{y \in W | \text{gr}(y) \leq u\}$.

Lemma 11.3. *If F_1, F_2 are free B_n -persistence modules with bases B_1, B_2 and $W_1 \subset F_1, W_2 \subset F_2$ are sets of homogeneous elements then a morphism $f : F_1 \rightarrow F_2$ maps $\langle W_1 \rangle$ into $\langle W_2 \rangle$ iff $[f, B_1, B_2][w, B_1] \in \text{span}[W_2^{gr(w)}, B_2]$ for every $w \in W_1$.*

Proof. This is straightforward. \square

11.2 Deciding Whether Two B_n -persistence Modules are ϵ -interleaved is Equivalent to Deciding the Solvability of a System in $\mathcal{MQ}(k)$

Let $\langle G_M | R_M \rangle, \langle G_N | R_N \rangle$ be presentations for finitely presented B_n -modules M and N , and assume the elements of each of the sets G_M, G_N, R_M, R_N are endowed with a total order, which may be chosen arbitrarily. For a finite ordered set T and $1 \leq i \leq |T|$, let $T_{,i}$ denote the i^{th} element of T .

We now define six matrices of variables, each with some of the variables constrained to be 0.

- Let \mathbf{A} be an $|G_N| \times |G_M|$ matrix of variables, with $\mathbf{A}_{ij} = 0$ iff $gr(G_{M,j}) < gr(G_{N,i}) + \epsilon$.
- Let \mathbf{B} be an $|G_M| \times |G_N|$ matrix of variables, with $\mathbf{B}_{ij} = 0$ iff $gr(G_{N,j}) < gr(G_{M,i}) + \epsilon$.
- Let \mathbf{C} be an $|R_N| \times |R_M|$ matrix of variables, with $\mathbf{C}_{ij} = 0$ iff $gr(R_{M,j}) < gr(R_{N,i}) + \epsilon$.
- Let \mathbf{D} be an $|R_M| \times |R_N|$ matrix of variables, with $\mathbf{D}_{ij} = 0$ iff $gr(R_{N,j}) < gr(R_{M,i}) + \epsilon$.
- Let \mathbf{E} be an $|R_M| \times |G_M|$ matrix of variables, with $\mathbf{E}_{ij} = 0$ iff $gr(G_{M,j}) < gr(R_{M,i}) + 2\epsilon$.
- Let \mathbf{F} be an $|R_N| \times |G_N|$ matrix of variables, with $\mathbf{F}_{ij} = 0$ iff $gr(G_{N,j}) < gr(R_{N,i}) + 2\epsilon$.

Let T_M denote the $|G_M| \times |R_M|$ matrix whose i^{th} column is $[R_{M,i}, G_M]$ and let T_N denote the $|G_N| \times |R_N|$ matrix whose i^{th} column is $[R_{N,i}, G_N]$.

Theorem 11.4. *M and N are ϵ -interleaved iff the multivariate system of quadratic equations*

$$\begin{aligned} \mathbf{A}T_M &= T_N\mathbf{C} \\ \mathbf{B}T_N &= T_M\mathbf{D} \\ \mathbf{B}\mathbf{A} - I_{|G_M|} &= T_M\mathbf{E} \\ \mathbf{A}\mathbf{B} - I_{|G_N|} &= T_N\mathbf{F} \end{aligned}$$

has a solution.

Proof. To prove the result, we proceed in three steps. First, we observe that for any free covers (F_M, ρ_M) and (F_N, ρ_N) of M and N , the existence of ϵ -interleaving morphisms between M and N is equivalent to the existence of a pair of morphisms between F_M and F_N having certain properties. We then note that the existence of such maps is equivalent to the existence of two matrices, one in $\text{Mat}_k(G_M, G_N)$ and the other in $\text{Mat}_k(G_N, G_M)$, having certain properties. Finally, we observe that the existence of such matrices is equivalent to the existence of a solution to the above multivariate system of quadratics.

Let (F_M, ρ_M) and (F_N, ρ_N) be free covers of M and N .

Lemma 11.5. *M and N are ϵ -interleaved iff there exist morphisms $\tilde{f} : F_M \rightarrow F_N(\epsilon)$ and $\tilde{g} : F_N \rightarrow F_M(\epsilon)$ such that*

1. $\tilde{f}(\ker(\rho_M)) \subset (\ker(\rho_N))(\epsilon)$,
2. $\tilde{g}(\ker(\rho_N)) \subset (\ker(\rho_M))(\epsilon)$,
3. $\tilde{g} \circ \tilde{f} - S(F_M, 2\epsilon) \subset (\ker(\rho_M))(2\epsilon)$,
4. $\tilde{f} \circ \tilde{g} - S(F_N, 2\epsilon) \subset (\ker(\rho_N))(2\epsilon)$.

We'll call morphisms \tilde{f}, \tilde{g} satisfying the above properties **ϵ -interleaved lifts** of the free covers (F_M, ρ_M) and (F_N, ρ_N) .

Proof. Let $f : M \rightarrow N(\epsilon)$ and $g : N \rightarrow M(\epsilon)$ be interleaving morphisms. Then by Lemma 2.1 there exist lifts $\tilde{f} : F_M \rightarrow F_N(\epsilon)$ and $\tilde{g} : F_N \rightarrow F_M(\epsilon)$ of f and g . By the definition of a lift, \tilde{f} and \tilde{g} satisfy properties 1 and 2 in the statement of the lemma. $\tilde{g} \circ \tilde{f}$ is a lift of $g \circ f = S(M, 2\epsilon)$. $S(F_M, 2\epsilon)$ is also a lift of $S(M, 2\epsilon)$, so by the uniqueness up to homotopy of lifts (Lemma 2.1), \tilde{f} and \tilde{g} satisfy property 3. The same argument shows that \tilde{f} and \tilde{g} satisfy property 4.

The converse direction is straightforward; we omit the details. \square

Now let $F_M = \langle G_M \rangle$, $F_N = \langle G_N \rangle$, and let $\rho_M : F_M \rightarrow F_M / \langle R_M \rangle$, $\rho_N : F_N \rightarrow F_N / \langle R_N \rangle$ be the quotient maps. Since the interleaving distance between two modules is an isomorphism invariant of the modules, we may assume without loss of generality that $F_M / \langle R_M \rangle = M$ and $F_N / \langle R_N \rangle = N$. Then (F_M, ρ_M) and (F_N, ρ_N) are free covers of M and N .

Lemma 11.6. *M and N are ϵ -interleaved iff there exist matrices $A \in \text{Mat}_k(G_M, G_N)$ and $B \in \text{Mat}_k(G_N, G_M)$ such that*

1. $A[w, G_M] \in \text{span}[R_N^{gr(w)+\epsilon}, G_N]$ for all $w \in R_M$,
2. $B[w, G_N] \in \text{span}[R_M^{gr(w)+\epsilon}, G_M]$ for all $w \in R_N$,
3. $(BA - I_{|G_M|})(\mathbf{e}_i) \in \text{span}[R_M^{gr(G_{M,i})+2\epsilon}, G_M]$ for $1 \leq i \leq |G_M|$,
4. $(AB - I_{|G_N|})(\mathbf{e}_i) \in \text{span}[R_N^{gr(G_{N,i})+2\epsilon}, G_N]$ for $1 \leq i \leq |G_N|$.

Proof. By Lemma 11.5, M and N are ϵ -interleaved iff there exists ϵ -interleaved lifts $\tilde{f} : F_M \rightarrow F_N$ and $\tilde{g} : F_N \rightarrow F_M$ of the free covers (F_M, ρ_M) and (F_N, ρ_N) .

By Lemma 11.3, morphisms $\tilde{f} : F_M \rightarrow F_N$ and $\tilde{g} : F_N \rightarrow F_M$, are ϵ -interleaved lifts iff

1. $[\tilde{f}, G_M, G_N(\epsilon)][w, G_M] \in \text{span}[R_N(\epsilon)^{gr(w)}, G_N(\epsilon)]$ for all $w \in R_M$,
2. $[\tilde{g}, G_N, G_M(\epsilon)][w, G_N] \in \text{span}[R_M(\epsilon)^{gr(w)}, G_M(\epsilon)]$ for all $w \in R_N$,
3. $[\tilde{g} \circ \tilde{f} - S(F_M, 2\epsilon), G_M, G_M(2\epsilon)][w, G_M] \in \text{span}[R_M(2\epsilon)^{gr(w)}, G_M(2\epsilon)]$ for all $w \in G_M$,
4. $[\tilde{f} \circ \tilde{g} - S(F_N, 2\epsilon), G_N, G_N(2\epsilon)][w, G_N] \in \text{span}[R_N(2\epsilon)^{gr(w)}, G_N(2\epsilon)]$ for all $w \in G_N$.

By Lemma 11.2,

$$[\tilde{g} \circ \tilde{f} - S(F_M, 2\epsilon), G_M, G_M(2\epsilon)] = [\tilde{g}, G_N, G_M(\epsilon)][\tilde{f}, G_M, G_N(\epsilon)] - I_{|G_M|}$$

and

$$[\tilde{f} \circ \tilde{g} - S(F_N, 2\epsilon), G_N, G_N(2\epsilon)] = [\tilde{f}, G_M, G_N(\epsilon)][\tilde{g}, G_N, G_M(\epsilon)] - I_{|G_N|}.$$

Also, by Remark 11.1, for $1 \leq i \leq |G_M|$, $[G_{M,i}, G_M] = \mathbf{e}_i$, where \mathbf{e}_i is the i^{th} standard basis vector in $k^{|G_M|}$. Similarly, for $1 \leq i \leq |G_N|$, $[G_{N,i}, G_N] = \mathbf{e}_i$, where \mathbf{e}_i is the i^{th} standard basis vector in $k^{|G_N|}$.

Finally, note that we have that

$$\begin{aligned} [R_N(\epsilon)^{gr(w)}, G_N(\epsilon)] &= [R_N^{gr(w)+\epsilon}, G_N] \text{ for all } w \in R_M, \\ [R_M(\epsilon)^{gw(w)}, G_M(\epsilon)] &= [R_M^{gr(w)+\epsilon}, G_M] \text{ for all } w \in R_N, \\ [R_M(2\epsilon)^{gr(w)}, G_M(2\epsilon)] &= [R_M^{gr(w)+2\epsilon}, G_M] \text{ for all } w \in G_M, \\ [R_N(2\epsilon)^{gr(w)}, G_N(2\epsilon)] &= [R_N^{gr(w)+2\epsilon}, G_N] \text{ for all } w \in G_N. \end{aligned}$$

Using all of these observations, Lemma 11.6 now follows from Lemma 11.1. \square

Finally, Theorem 11.4 follows from Lemma 11.6 by way of elementary matrix algebra and, in particular, the basic fact that for $l, m \in \mathbb{N}$ and vectors v, v_1, \dots, v_l in k^m , $v \in \text{span}(v_1, \dots, v_l)$ iff there exists a vector $w \in k^l$ such that $v = Vw$, where V is the $m \times l$ matrix whose i^{th} column is v_i . \square

Remark 11.2. Note that the size of the system of quadratic equations in the statement of Theorem 11.4 is $O(q^2)$, where q is the total number of generators and relations in the presentations for M and N . For any $\epsilon \geq 0$, the system of quadratics has as few variables and equations as possible when the presentations for M and N are minimal.

11.3 Determining Possible Values for $d_I(M, N)$

Let M and N be finitely presented B_n -modules, and let U_M^i, U_N^i, U_M , and U_N be as defined at the beginning of Section 6. Let

$$U_{M,N} = \bigcup_i \left(\{|x - y|\}_{x \in U_M^i, y \in U_N^i} \cup \left\{ \frac{1}{2}|x - y| \right\}_{x, y \in U_M^i} \cup \left\{ \frac{1}{2}|x - y| \right\}_{x, y \in U_N^i} \right) \cup \{0, \infty\}.$$

Note that $|U_{M,N}| = O(q^2)$, where as above q is the total number of generators and relations in a minimal presentation for M and a minimal presentation for N .

Proposition 11.7. $d_I(M, N) \in U_{M,N}$.

Proof. Assume that for some $\epsilon' > 0$, $\epsilon' \notin U_{M,N}$, M and N are ϵ' -interleaved. Let ϵ be the largest element of $U_{M,N}$ such that $\epsilon' > \epsilon$, and let $\delta = \epsilon' - \epsilon$.

We'll check that M, N, ϵ and δ satisfy the hypotheses of Lemma 6.7. The lemma then implies that M and N are ϵ -interleaved. The result follows.

By assumption, M and N are $(\epsilon + \delta)$ -interleaved, so the first hypothesis of Lemma 6.7 is satisfied. We'll show that the second hypothesis is satisfied; the proof that the third hypothesis is satisfied is the same as that for the second hypothesis.

If $z \in U_M$ then for no i , $1 \leq i \leq n$, can an element of U_N^i lie in $(z_i + \epsilon, z_i + \epsilon + \delta]$; if, to the contrary, for some i there were an element $u \in U_N^i \cap (z + \epsilon, z + \epsilon + \delta]$, then we would have $|u - z_i| \in U_{M,N}$, and $\epsilon < |u - z_i| \leq \epsilon + \delta$, which contradicts the way we chose ϵ and δ . Thus by Lemma 6.4, $\varphi_N(z + \epsilon, z + \epsilon + \delta)$ is an isomorphism.

Similarly, for no i , $1 \leq i \leq n$, can an element of U_M^i lie in $(z_i + 2\epsilon, z_i + 2\epsilon + 2\delta]$; if, to the contrary, for some i there were an element $u \in U_M^i \cap (z + 2\epsilon, z + 2\epsilon + 2\delta]$, then we would have

$\frac{1}{2}|u - z_i| \in U_{M,N}$, and $\epsilon < \frac{1}{2}|u - z_i| \leq \epsilon + \delta$, which again contradicts the way we chose ϵ and δ . By Lemma 6.4, $\varphi_M(z + 2\epsilon, z + 2\epsilon + 2\delta)$ is an isomorphism.

Thus the second hypothesis of Lemma 6.7 is satisfied by our M, N, ϵ, δ , as we wanted to show. \square

12 Discussion of Future Work

We believe that Theorem 5.2, Corollary 6.2, and Corollary 10.2 establish the credentials of the interleaving distance as a natural generalization of the bottleneck distance to the setting of multidimensional persistence.

Insofar as the interleaving distance is in fact a good choice of distance on multidimensional persistence modules, the question of how to compute it is interesting and, it seems to us, potentially important from the standpoint of applications. The results of Section 11 suggest a path towards the development of a theory of computation of the interleaving distance. We plan to pursue this path further in subsequent work. As noted in Section 11, to exploit the connection with multivariate quadratics in the development of such a theory in practice, one needs in particular a way of computing minimal presentations of simplicial homology modules of simplicial n -filtrations. We will address this problem in a companion paper.

As mentioned in the introduction, our view is that the existence of a good choice of metric on multidimensional persistence modules promises to facilitate the adaptation to the multidimensional setting of theoretical results and applications of ordinary persistence which depend on the bottleneck distance. We believe that there is plenty of interesting and potentially useful work to be done in carrying out this adaptation.

Corollary 10.2 demonstrates that the interleaving distance is optimal in the sense of Example 9.1 when $k = \mathbb{Q}$ or $\mathbb{Z}/p\mathbb{Z}$. However, our discussion of optimality of pseudometrics in Section 9 raises many more questions than it answers. Some of the more interesting questions are:

1. Can we extend the result of Theorem 10.1 to arbitrary ground fields? It seems this would involve invoking an analogue of the universal coefficient theorem for persistence.
2. Can we extend the result of Theorem 10.1 to the case $i = 0$?
3. Can we prove that the interleaving metric is \mathcal{R} -optimal for \mathcal{R} any of the relative structures on $\text{obj}^*(B_n\text{-mod})$ defined in Examples 9.3-9.5? The case of Example 9.4 is of particular interest to us. We have observed in Section 9.3 that in this case an \mathcal{R} -optimal metric does exist.
4. Can we obtain analogous results about the optimality of metrics on more general types of persistent homology modules? For instance, can we prove a result analogous to Theorem 10.1 for levelset zigzag persistence [5]?

An interesting question related to question 4 above is whether there is a way of algebraically reformulating the bottleneck distance for zigzag persistence modules as an analogue of the interleaving distance in such way that the definition generalizes to a larger classes of quiver representations [21].

Finally, we mention again that it would be nice to have an extension of Theorem 4.5 to a structure theorem for arbitrary tame B_1 -persistence modules, and an extension of Corollary 6.2 to well behaved tame B_n -persistence modules.

Acknowledgments

My discussions with my adviser Gunnar Carlsson have catalyzed the research presented here in several key ways. In our many conversations about multidimensional persistent homology, Gunnar has impressed on me the value of bringing the machinery of commutative algebra to bear on the study of multidimensional persistence. At the same time, he has served as a patient and helpful sounding board for the ideas of this paper. I thank Gunnar for his support and guidance, for suggesting the use of the theory of coherent rings in this paper, and for valuable input on how to cleanly explain the functorial relationship between $k[G_n]$ -persistence modules and \mathbf{A}_n -persistence modules.

This work was supported by a research assistantship funded through ONR grant number N00014-09-1-0783.

References

- [1] S. Biasotti, A. Cerri, P. Frosini, D. Giorgi, and C. Landi. Multidimensional size functions for shape comparison. *Journal of Mathematical Imaging and Vision*, 32(2):161–179, 2008.
- [2] G.E. Bredon. *Topology and geometry*. Springer Verlag, 1993.
- [3] E. Carlsson, G. Carlsson, V. De Silva, and SJ Fortune. An algebraic topological method for feature identification. *International Journal of Computational Geometry and Applications*, 16(4):291–314, 2006.
- [4] G. Carlsson. Topology and data. *American Mathematical Society*, 46(2):255–308, 2009.
- [5] G. Carlsson, V. De Silva, and D. Morozov. Zigzag persistent homology and real-valued functions. In *Proceedings of the 25th annual symposium on Computational geometry*, pages 247–256. ACM, 2009.
- [6] G. Carlsson, T. Ishkhanov, V. De Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *International Journal of Computer Vision*, 76(1):1–12, 2008.
- [7] G. Carlsson, G. Singh, and A. Zomorodian. Computing multidimensional persistence. *Algorithms and Computation*, pages 730–739, 2009.
- [8] G. Carlsson and A. Zomorodian. The theory of multidimensional persistence. *Discrete and Computational Geometry*, 42(1):71–93, 2009.
- [9] G. Carlsson, A. Zomorodian, A. Collins, and L. Guibas. Persistence barcodes for shapes. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 124–135. ACM, 2004.
- [10] A. Cerri, B. Di Fabio, M. Ferri, P. Frosini, and C. Landi. Multidimensional persistent homology is stable. *Arxiv preprint arXiv:0908.0064*, 2009.

- [11] F. Chazal, D. Cohen-Steiner, M. Glisse, L.J. Guibas, and S.Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the 25th annual symposium on Computational geometry*, pages 237–246. ACM, 2009.
- [12] F. Chazal, D. Cohen-Steiner, L.J. Guibas, F. Mémoli, and S.Y. Oudot. Gromov-Hausdorff stable signatures for shapes using persistence. In *Proceedings of the Symposium on Geometry Processing*, pages 1393–1403. Eurographics Association, 2009.
- [13] F. Chazal, L.J. Guibas, S. Oudot, and P. Skraba. Persistence-based clustering in Riemannian manifolds. 2009.
- [14] F. Chazal, L.J. Guibas, S.Y. Oudot, and P. Skraba. Analysis of scalar fields over point cloud data. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1021–1030. Society for Industrial and Applied Mathematics, 2009.
- [15] F. Chazal and S.Y. Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, pages 232–241. ACM, 2008.
- [16] C. Chen and D. Freedman. Measuring and computing natural generators for homology groups. *Computational Geometry*, 43(2):169–181, 2010.
- [17] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete and Computational Geometry*, 37(1):103–120, 2007.
- [18] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and Y. Mileyko. Lipschitz functions have L_p -stable persistence. *Foundations of Computational Mathematics*, 10(2):127–139, 2010.
- [19] M. d’Amico, P. Frosini, and C. Landi. Natural pseudo-distance and optimal matching between reduced size functions. *Acta applicandae mathematicae*, 109(2):527–554, 2010.
- [20] V. de Silva and R. Ghrist. Coverage in sensor networks via persistent homology. *Algebraic and Geometric Topology*, 7:339–358, 2007.
- [21] H. Derksen and J. Weyman. Quiver representations. *Notices of the AMS*, 52(2):200–206, 2005.
- [22] D.S. Dummit and R.M. Foote. *Abstract algebra*. Wiley, 1999.
- [23] H. Edelsbrunner and J. Harer. Persistent homology: a survey. In *Surveys on discrete and computational geometry: twenty years later: AMS-IMS-SIAM Joint Summer Research Conference, June 18-22, 2006, Snowbird, Utah*, volume 453, page 257. Amer Mathematical Society, 2008.
- [24] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28(4):511–533, 2002.
- [25] D. Eisenbud. *Commutative algebra with a view toward algebraic geometry*. Springer, 1995.
- [26] S. Glaz. *Commutative coherent rings*. Springer, 1989.
- [27] A. Hatcher. *Algebraic topology*. Cambridge Univ Pr, 2002.
- [28] S. Mac Lane. *Categories for the working mathematician*, volume 5. Springer verlag, 1998.

- [29] N. Milosavljevic, D. Morozov, and P. Skraba. Zigzag Persistent Homology in Matrix Multiplication Time. *Proceedings of the Annual Symposium on Computational Geometry*, 2011.
- [30] P. Skraba, M. Ovsjanikov, F. Chazal, and L. Guibas. Persistence-based segmentation of deformable shapes. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*, pages 45–52. IEEE, 2010.
- [31] A. Verri, C. Uras, P. Frosini, and M. Ferri. On the use of size functions for shape analysis. *Biological Cybernetics*, 70(2):99–107, 1993.
- [32] C. Webb. Decomposition of graded modules. *AMERICAN MATHEMATICAL SOCIETY*, 94(4), 1985.
- [33] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete and Computational Geometry*, 33(2):249–274, 2005.

A Appendix: The Coherence of B_n

A.1 Coherence: Basic Definitions and Results

$k[x_1, \dots, x_n]$ is well known to be a Noetherian ring. Finitely generated modules over Noetherian rings have some very nice algebraic properties. Here we define a standard weakening of the Noetherian property called coherence. Analogues of many of the same nice algebraic properties that hold for finitely generated modules over Noetherian rings hold for finitely presented modules over coherent rings. In particular, we have Corollary A.4, which we will use in Appendix B to prove Theorem 2.2.

Definition. For R a ring, we say an R -module M is **coherent** if M is finitely generated and every finitely generated submodule of M is finitely presented. We say a ring R is coherent if it is a coherent module over itself.

Coherent commutative rings and coherent modules are well studied; the following results are standard. The reader may refer to [26] for the proofs.

Proposition A.1. *If R is a Noetherian ring then R is coherent.*

Theorem A.2. *If R is a coherent ring then every finitely presented R -module is coherent.*

Theorem A.3. *If $f : M \rightarrow N$ is a morphism between coherent R -modules M and N then $\ker(f)$, $\operatorname{im}(f)$, and $\operatorname{coker}(f)$ are coherent R -modules.*

Combining these last two theorems immediately gives

Corollary A.4. *If R is a coherent ring and $f : M \rightarrow N$ is a morphism between finitely presented R -modules M and N , then $\ker(f)$, $\operatorname{im}(f)$, and $\operatorname{coker}(f)$ are finitely presented.*

A.2 The Ring B_n is Coherent

Theorem A.5. *For any $n \in \mathbb{N}$, B_n is coherent.*

Proof. The key to the proof is the following theorem:

Theorem A.6 ([26, Theorem 2.3.3]). *Let $\{R_\alpha\}_{\alpha \in S}$ be a directed system of rings and let $R = \lim_{\rightarrow} R_\alpha$. Suppose that for $\alpha \leq \beta$, R_β is a flat R_α module and that R_α is coherent for every α . Then R is a coherent ring.*

First, recall that \mathbb{R} is a vector space over \mathbb{Q} . We'll say that $a_1, \dots, a_l \in \mathbb{R}_{\geq 0}$ are *rationally independent* if they are linearly independent as vectors in \mathbb{R} over the field \mathbb{Q} .

We next extend this definition to vectors in $\mathbb{R}_{\geq 0}^n$: We say a finite set $V \subset \mathbb{R}_{\geq 0}^n$ is *rationally independent* if

1. V is the union of sets V_1, \dots, V_n , where each element of V_i has a non-zero i^{th} coordinate and all other coordinates are equal to zero.
2. For any i , if a_1, \dots, a_l are the non-zero coordinates of the elements of V_i (listed with multiplicity), then a_1, \dots, a_l are rationally independent in the sense defined above.

We define an n -grid to be a monoid generated by some rationally independent set $V \subset \mathbb{R}_{\geq 0}^n$. Denote the n -grid generated by the rationally independent set V as $\Gamma(V)$. $\Gamma(V)$ is a submonoid of $\mathbb{R}_{\geq 0}^n$.

Lemma A.7. *If V is a rationally independent set, then the n -grid generated by V is isomorphic to $\mathbb{Z}_{\geq 0}^{|V|}$.*

Proof. The proof is straightforward; we omit it. □

As noted in Section 2.2, for any $m \in \mathbb{N}$, $k[\mathbb{Z}_{\geq 0}^m] \cong k[x_1, \dots, x_m]$. As the latter ring is Noetherian, it is coherent by proposition A.1. Thus if G is an n -grid, $k[G]$ is coherent.

Lemma A.8. *For any finite set $A \subset \mathbb{R}_{\geq 0}$, there's a rationally independent set $B \subset \mathbb{R}_{\geq 0}$ such that A lies in the monoid generated by B .*

Proof. We proceed by induction on the number of elements l in the set A . The base case is trivial. Now assume the result holds for sets of order $l - 1$. Write $A = \{a_1, \dots, a_l\}$. By the induction hypothesis there exists a finite rationally independent set $A' = \{a'_1, \dots, a'_m\}$ such that $\{a_1, \dots, a_{l-1}\}$ lies in $\Gamma(A')$. If $A' \cup a_l$ is rationally independent, take $B = A' \cup a_l$. Otherwise $a_l = q_1 a'_1 + \dots + q_{m-1} a'_{m-1}$ for some $q_1, \dots, q_{m-1} \in \mathbb{Q}$; we may take $B = \{q'_1 a'_1, \dots, q'_{l-1} a'_{m-1}\}$, where $q'_i = 1/b_i$ for some $b_i \in \mathbb{N}$ such that $q_i = a/b_i$ for some $a \in \mathbb{Z}_{\geq 0}$. □

Lemma A.9. *The set of n -grids forms a directed system under inclusion with direct limit $\mathbb{R}_{\geq 0}^n$.*

Proof. To show that the set of n -grids forms a directed system, we need that given two n -grids G_1 and G_2 , there's an n -grid G_3 such that $G_1 \subset G_3$ and $G_2 \subset G_3$. This follows readily from Lemma A.8; we leave the details to the reader. Any element of $\mathbb{R}_{\geq 0}^n$ lies in an n -grid, so $\mathbb{R}_{\geq 0}^n$ must be the colimit of the directed system. □

For a monoid A and a submonoid $A' \subset A$, we have $k[A'] \subset k[A]$. This implies the following:

Lemma A.10. *The set of rings $\{k[G] \mid G \text{ is an } n\text{-grid}\}$ has the structure of a directed system induced by the directed system structure on the set of n -grids, and B_n is the direct limit of this directed system.*

Proposition A.11. *Given two positive n -grids G', G with $G' \subset G$, $k[G]$ is a free $k[G']$ module.*

Proof. We begin by establishing a couple of lemmas.

Lemma A.12. *For any rationally independent set V' and n -grid A containing $\Gamma(V')$, there is a rationally independent set V such that $A = \Gamma(V)$ and such that for each $a \in V'$, V contains an element of the form a/b for some $b \in \mathbb{N}$.*

We call V an *extension* of V' .

Proof. The proof of Lemma A.12 is similar to the proof of Lemma A.8; we omit it. \square

Let \mathcal{S} denote the set of maximal sets of the form $g + G' \equiv \{g + g' \mid g' \in G'\}$ for some $g \in G$.

Lemma A.13. *The sets \mathcal{S} form a partition of G .*

Proof. It's enough to show that if $g_1 + G', g_2 + G' \in \mathcal{S}$ and $g_1 + G' \cap g_2 + G' \neq \emptyset$, then $g_1 + G' = g_2 + G'$.

Let V' be a rationally independent set with $\Gamma(V') = G'$, and let V be an extension of V' with $\Gamma(V) = G$. Write $V' = \{v_1, \dots, v_l\}$ and $V = \{v_1/b_1, \dots, v_l/b_l, v_{l+1}, \dots, v_m\}$ for some $b_1, \dots, b_l \in \mathbb{N}$.

Assume there exist $g'_1, g'_2 \in G'$ such that $g_1 + g'_1 = g_2 + g'_2$. We'll show that there then exists an element $g_3 \in G$ such that $g_1, g_2 \in g_3 + G$. By the maximality of $g_1 + G'$ and $g_2 + G'$, this implies $g_1 + G' = g_2 + G'$, as needed. We write

$$\begin{aligned} g_1 &= y_1 v_1/b_1 + \dots + y_l v_l/b_l + y_{l+1} v_{l+1} + \dots + y_m v_m, \\ g_2 &= z_1 v_1/b_1 + \dots + z_l v_l/b_l + z_{l+1} v_{l+1} + \dots + z_m v_m, \\ g'_1 &= y'_1 v_1 + \dots + y'_l v_l, \\ g'_2 &= z'_1 v_1 + \dots + z'_l v_l. \end{aligned}$$

for some $y_1, \dots, y_m, z_1, \dots, z_m, y'_1, \dots, y'_l, z'_1, \dots, z'_l \in \mathbb{Z}$. By the rational independence of V and the fact that $g_1 + g'_1 = g_2 + g'_2$, we have that $y_i = z_i$ for $l+1 \leq i \leq m$.

Define

$$\begin{aligned} g_3 &= \min(y_1, z_1) v_1/b_1 + \dots + \min(y_l, z_l) v_l/b_l + y_{l+1} v_{l+1} + \dots + y_m v_m, \\ g''_1 &= (y_1 - \min(y_1, z_1)) v_1/b_1 + \dots + (y_l - \min(y_l, z_l)) v_l/b_l, \\ g''_2 &= (z_1 - \min(y_1, z_1)) v_1/b_1 + \dots + (z_l - \min(y_l, z_l)) v_l/b_l. \end{aligned}$$

$g_3 + g''_1 = g_1$ and $g_3 + g''_2 = g_2$, so if we can show that $g''_1, g''_2 \in G'$ we are done.

By the rational independence of V and the fact that $g_1 + g'_1 = g_2 + g'_2$, for $1 \leq i \leq l$ we have that $\min(y_i, z_i)/b_i + \max(y'_i, z'_i) = y_i/b_i + y'_i$. This implies that $\max(y'_i, z'_i) - y'_i = (y_i - \min(y_i, z_i))/b_i$. In particular, the term on the right hand side lies in $\mathbb{Z}_{\geq 0}$. Thus $g''_1 \in G'$. The same argument shows $g''_2 \in G'$. \square

Now we are ready to complete the proof of Proposition A.11. It's easy to see that for any $s \in \mathcal{S}$, the natural action of G' on s extends to give $k[s]$ the structure of a free $k[G']$ module of rank 1. It follows from Lemma A.13 that the sets $\{k[s]\}_{s \in \mathcal{S}}$ have trivial intersection as $k[G']$ -submodules of $k[G]$. We then have that as a $k[G']$ module, $k[G] = \bigoplus_{s \in \mathcal{S}} k[s]$, and so in particular $k[G]$ is a free $k[G']$ -module, as we wanted to show. \square

Given Lemma A.10 and Proposition A.11, Theorem A.6 applies to give that B_n is coherent, since free modules are flat [25]. \square

B Minimal Presentations of B_n -persistence Modules

This section is devoted to the proof of Theorem 2.2.

B.1 Free Hulls

We first observe that some standard results about resolutions and minimal resolutions of modules over local rings adapt to B_n -persistence modules. We'll only be interested in the specialization of such results to the 0^{th} modules in a free resolution, and for the sake of simplicity we phrase the results only for this special case. However, the results discussed here do extend to statements about free resolutions of finitely presented B_n -persistence modules.

Let \mathfrak{m} denote the ideal of B_n generated by the set

$$\{v \in B_n \mid v \text{ is homogeneous and } gr(v) > 0\}.$$

Define a *free hull* of M to be a free cover (F_M, ρ_M) such that $\ker(\rho_M) \subset \mathfrak{m}F_M$.

Nakayama's lemma [25] is a key ingredient in the proofs of the results about free resolutions over local rings that we would like to adapt to our setting. To adapt these proofs, we need an n -graded version of Nakayama's lemma.

Lemma B.1 (Nakayama's Lemma for Persistence Modules). *Let M be a finitely generated B_n -persistence module. If $y_1, \dots, y_m \in M$ have images in $M/\mathfrak{m}M$ that generate the quotient, then y_1, \dots, y_m generate M .*

Proof. The usual Proof of Nakayama's lemma [25] carries over with only minor changes. \square

Lemma B.2. *A free cover (F_M, ρ_M) of a finitely generated B_n -persistence module M is a free hull iff a basis for F_M maps under ρ_M to a minimal set of generators for M .*

Proof. Given the adaptation Lemma B.1 of Nakayama's lemma to our setting, the proof of [25, Lemma 19.4] gives the result. \square

It follows easily from Lemma B.2 that a free hull exists for any finitely generated B_n -persistence module M . Corollary B.4 below gives a uniqueness result for free hulls.

Theorem B.3. *If (F_M, ρ_M) is a free hull of a finitely presented B_n -persistence module M and (F'_M, ρ'_M) is any free cover of M , then F_M includes as a direct summand of F'_M in such a way that $F'_M \cong F_M \oplus F''_M$ for some free module F''_M , and $\ker(\rho'_M) = \ker(\rho_M) \oplus F''_M \subset F_M \oplus F''_M$.*

Sketch of Proof. The statement of the theorem is the specialization to 0^{th} modules in the free resolutions of M of an adaptation of [25, Theorem 20.2] to our B_n -persistence setting. To modify Eisenbud's proof of [25, Theorem 20.2] to obtain a proof of Theorem B.3, one needs to invoke the coherence of B_n and use Corollary A.4 to show that $\ker(\rho_M)$ is finitely generated. Given this, the strategy of proof adapts in a straightforward way. \square

Corollary B.4 (Uniqueness of free hulls). *If M is a finitely presented B_n -persistence module, and (F_M, ρ_M) , (F'_M, ρ'_M) are two free hulls of M , then there is an isomorphism from F_M to F'_M which is a lift of the identity map of M .*

Proof. By Theorem B.3, we can identify F_M with a submodule of F'_M in such a way that $F'_M = F_M \oplus F''_M$ for some free module F''_M and $\ker(\rho'_M) = \ker(\rho) \oplus F''_M \subset F_M \oplus F''_M$. Since F'_M is a free hull, we must have $\ker(\rho'_M) \in \mathfrak{m}F'_M$, which implies $F''_M = 0$. The result follows. \square

Corollary B.5. *If M is a finitely presented B_n -persistence module and B, B' are two minimal sets of generators for M , then $gr(B) = gr(B')$.*

Proof. This follows from Corollary B.4 and Lemma B.2. \square

B.2 Proof of Theorem 2.2

Recall that a minimal presentation $\langle G|R \rangle$ of a B_n -persistence module M is one such that

1. the quotient $\langle G \rangle \rightarrow \langle G \rangle / \langle R \rangle$ maps G to a minimal set of generators for $\langle G \rangle / \langle R \rangle$.
2. R is a minimal set of generators for $\langle R \rangle$.

Let M be a finitely presented B_n -persistence module. Let $\langle G|R \rangle$ be a minimal presentation of M . We need to show that for any other presentation $\langle G'|R' \rangle$ of M , $gr(G) \leq gr(G')$ and $gr(R) \leq gr(R')$.

Let $\psi : \langle G \rangle / \langle R \rangle \rightarrow M$ and $\psi' : \langle G' \rangle / \langle R' \rangle \rightarrow M$ be isomorphisms, let $\pi : \langle G \rangle \rightarrow \langle G \rangle / \langle R \rangle$ and $\pi' : \langle G' \rangle \rightarrow \langle G' \rangle / \langle R' \rangle$ be the quotient homomorphisms, let $\rho = \psi \circ \pi$, and let $\rho' = \psi' \circ \pi'$. Then by Lemma B.2, $(\langle G \rangle, \rho)$ is a free hull of M , and $(\langle G' \rangle, \rho')$ is a free cover of M .

By Theorem B.3, $\langle G \rangle$ includes as a direct summand of $\langle G' \rangle$. The image of G under this inclusion can be extended to a basis for $\langle G' \rangle$. Recall that if B and B' are two bases for a free B_n -persistence module F , then $gr(B) = gr(B')$. We thus have that $gr(G) \leq gr(G')$.

Theorem B.3 also implies that $\langle R' \rangle \cong \langle R \rangle \oplus F$ for some free B_n -persistence module F . Let B be a basis for F . Then $R \cup B$ is a minimal set of generators for $\langle R \rangle \oplus F$. Let R'' denote the image of R' under an isomorphism from $\langle R' \rangle$ to $\langle R \rangle \oplus F$ and let $p : \langle R \rangle \oplus F \rightarrow \langle R \rangle$ denote projection onto the first summand. Since p is surjective, $p(R'')$ is a set of homogeneous generators for $\langle R \rangle$.

Since $\langle G \rangle$ and M are finitely presented, by Corollary A.4 $\ker(\rho) = \langle R \rangle$ is also finitely presented. Then by Corollary B.5, $gr(R) \leq gr(p(R''))$. Since $gr(p(R'')) \leq gr(R'') = gr(R')$ we have that $gr(R) \leq gr(R')$. \square